

The Macroeconomic Effects of Neighborhood Policies: A Dynamic Analysis*

Alessandra Fogli [†] Veronica Guerrieri [‡] Mark Ponder [§] Marta Prato [¶]

April 9, 2026

Abstract

We study the macroeconomic effects of neighborhood-specific policies in a general equilibrium model of a city with endogenous residential sorting and educational investment. A key feature of the model is the presence of endogenous local spillovers that depend on the distribution of families across neighborhoods. We analyze three policies: a housing-voucher policy inspired by the MTO program, which enables poor families to relocate to low-poverty neighborhoods; a place-based transfer (PBT) policy that provides monetary transfers to families in poor neighborhoods; and a place-based investment (PBI) policy that invests resources in local institutions, such as public schools, to directly enhance local spillovers. We find that the MTO policy generates substantial welfare gains for recipient families, but scaling up the program dampens these gains and induces large welfare losses for non-recipients. By contrast, the PBT policy delivers larger average welfare gains but is less effective in reducing inequality and segregation. Finally, the PBI policy produces smaller short-run effects but, over time, resolves the trade-off by raising average welfare while simultaneously reducing inequality, lowering segregation, and improving intergenerational mobility.

Keywords: Income inequality, Residential segregation, MTO, Place-Based policies, Local externalities
JEL Classification: E24, I2, O15, R2

*For helpful comments, we are grateful to our editors Gianluca Violante and Valerie Ramey. We also thank Francesco Agostinelli, Humlum Anders, Adrien Auclert, Job Boerma, Diego Daruich, Jesse Gregory, Kyle Herkenhoff, Erik Hurst, Rishabh Kirpalani, Paolo Martellini, Simon Mongey, and Ludwig Straub, and several seminar participants. For outstanding research assistance, we thank Dhananjay Ghei and Bipul Verma. A preliminary version of this paper has circulated with the title “Scaling Up the American Dream: A Dynamic Analysis”. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Minneapolis, the Federal Reserve System.

[†]Federal Reserve Bank of Minneapolis, afogli00@gmail.com

[‡]University of Chicago, Booth School of Business, vguerrie@chicagobooth.edu

[§]NERA, ponder.mark@gmail.com

[¶]Bocconi University, IGIER, and CEPR, marta.prato@unibocconi.it

1 Introduction

In recent decades, the United States has witnessed a marked rise in income inequality and residential segregation, trends that have profound implications for social mobility and access to opportunity. As income disparities have widened, so too have the geographical divides between affluent and disadvantaged communities. This increasing concentration of wealth and poverty has sparked ongoing debates over the efficacy of neighborhood-specific policies aimed at addressing these inequalities. These discussions are crucial in understanding how targeted interventions can reshape the fabric of our cities and foster greater economic inclusion.

A key example is the Moving to Opportunity (MTO) program, launched in the mid-1990s, which provided housing vouchers to low-income families in high-poverty neighborhoods. Evidence from Chetty et al. (2016) shows that relocating to lower-poverty areas significantly improved children's long-term outcomes, particularly in education and earnings.

Other place-based policies, such as the Earned Income Tax Credit (EITC), also play a redistributive role. The EITC raises after-tax income for low-wage households, with especially strong effects in distressed neighborhoods. Alongside housing subsidies and food assistance, these transfers reduce financial hardship and support household well-being.

Policymakers have also worked to reduce opportunity gaps through neighborhood-focused investments aimed at improving both education and infrastructure. One key policy in this regard is Title I funding, which allocates additional federal resources to schools in low-income areas to support students' academic achievement. Another example is the Community Development Block Grants (CDBG) program, which channels federal funds into neighborhood revitalization projects. Together, these policies seek to improve local living conditions and expand pathways to economic mobility.

These different types of neighborhood-specific policies have different effects on the sorting of families across neighborhoods and, hence, on aggregate outcomes. In this paper, we use a general equilibrium model with endogenous local spillovers, residential choice and educational investment, to compare the macroeconomic impact of these policies over scale and over time. In particular, local spillovers capture different channels through which neighborhoods affect children's outcomes: from the quality of public schools, to peer effects, social norms, networks, and so forth.

We consider a model of a representative city with three neighborhoods, inhabited by overlapping generations of agents that live for two periods. In the first period, an agent is a child, and in the second, she becomes a parent. Parents care about their own consumption and their children's expected future income, and choose how much to consume, which neighborhood to raise their children in, and how much to invest in their children's education. Children have heterogeneous ability and their future wages are affected by the neighborhood where they grow up because of the local spillover. In particular, we assume that the local spillover depends on the average expected income of the children growing up in that neighborhood, which means that the spillover tends to be higher in neighborhoods with richer parents and with higher ability

children. We assume that the spillover is complementary to both education and ability, so that living in a neighborhood with higher spillover generates higher returns to education, especially for the more talented children. The presence of the local spillovers generates residential sorting, as richer parents tend to choose to live in the neighborhood with the highest spillover to give their kids the best opportunity for success. We also assume that there is a government, which collects proportional income taxes to finance a public assistance program that provides free housing and a minimum level of consumption for poor families.

We calibrate the model to a representative US metro area in 1980 and then assume that in 1990 the economy is hit by a skill premium shock, that increases inequality, and, in turn, segregation. To discipline the strength of the local spillover, we use the micro estimates by Chetty and Hendren (2018) for the local exposure effects on children's outcomes. We then compare the effects of the introduction of three types of unexpected and permanent neighborhood-specific policies: a housing voucher scheme that mirrors the MTO program, a place-based transfer policy (PBT), and a place-based investment policy (PBI). We first compare the impact effects of the three policies across different financing scales, and then we compare the dynamic implications for a given scale.

We start by exploring the effects of the MTO policy, which offers a housing voucher to poor families living in the poorest neighborhood to move to the richest one. We first focus on a small financing scale that matches the effective MTO program and show that the model is able to generate a children's income gain for recipient families that is close to the estimate in Chetty et al. (2016) and a take-up rate close to the data. At this small scale, general equilibrium effects are negligible. However, we then explore the effects of scaling up the policy and show that general equilibrium effects become sizeable: as a larger fraction of poor families living in the poorest neighborhood receive a voucher to move to the richest neighborhood, the size of the spillover in the three neighborhoods endogenously change, as they depend on the distribution of families living there. In particular, the size of the spillover in the richest neighborhood decreases with the scale of the program. This dampens the income gain for the children of voucher-recipient families and generates income losses for children of families who used to live in that neighborhood. This implies that the MTO policy, while successful in decreasing inequality and segregation, may have small or even negative effects on average utilitarian welfare depending on the scale.

We then move to analyzing the effects of the PBT policy for the same financing scales. This policy offers a monetary transfer to all families living in the poorest neighborhood. This implies that the size of the transfer going to each family is much smaller than the voucher under the MTO policy, generating smaller welfare gains for the recipient families. However, the general equilibrium effects are drastically different. In particular, the PBT policy incentivizes families to move to the neighborhood with the lowest spillover to receive the transfer. This implies that the poorest families living in the richest neighborhood move out, hence making that neighborhood even richer. This feeds back into an even higher spillover in the richest neighborhood. It follows that under PBT policy, both non-recipient and recipient families experience welfare gains. Overall the average city-wide welfare gains end up being larger than under the MTO policy, even though the welfare gains for the recipient families are smaller. At the same time, as the general equilibrium

effects generate more sorting by income across neighborhoods, the PBT policy is less successful in reducing inequality and segregation.

Finally, we explore the effects of the PBI policy that uses the same finances to invest directly in increasing the size of the local spillover (for example, by investing in public schools) and financing a minimum level of education for the poorer families. To discipline the design of this policy, we use the estimates about return from capital investment in public schools from Biasi et al. (2025). Such a policy is not as successful as the others in improving city-wide welfare on impact. However, the effect of the spillover accumulates over time and generates larger welfare gain in the long run, and, at the same time, succeeds in reducing both inequality and residential segregation by income and in improving intergenerational mobility.

Related Literature

Our model builds on the theoretical literature that studies inequality and residential segregation in general equilibrium models with local spillovers and residential choice, going back to Benabou (1993), Benabou (1996a), Benabou (1996b), Durlauf (1996a), Durlauf (1996b), Fernandez and Rogerson (1996), and Fernandez and Rogerson (1998).

There has been a recent growing body of work using these types of models to study the effects of different neighborhood-specific policies.¹ We build on the model in Fogli et al. (2026b) that proposes a quantitative GE model with local spillovers to study the role of residential segregation in amplifying the response of inequality to skill premium shock. The paper that is closest to ours is Chyn and Daruich (2022), which also explores the effects of housing vouchers and place-based policies in a general equilibrium model with endogenous neighborhood externalities. Our papers are complementary. Their model is richer in the description of the skill-accumulation technology and focuses on how endogenous labor supply responds to policies and how this affects parents' investment in child development. We model parents' investment in children's education in terms of economic resources rather than time. Given the different modeling choices, we analyze different types of place-based policies, with particular focus on the effects on residential segregation by income. Moreover, our model features three neighborhoods instead of two, which allows us to generate richer moving patterns, as families living in the policy destination neighborhoods may decide to move out into middle-income areas, reducing the welfare gains from the MTO policy that one would obtain with only two neighborhoods.² This is an important mechanism, given recent findings in Derenoncourt (2022) analyzing the Great Migration of Black families from the South to Northern cities between 1940 and 1960. She shows that the negligible effects on migrants' socio-economic outcomes were not due to negative selection in migration, but rather to the endogenous response of white residents, who moved out of the destination neighborhoods, undermining the potential gains from relocation. This mechanism is closely related to

¹ See Fogli et al. (2026a) for an overview of this class of models.

² To highlight the importance of three neighborhoods, in Appendix D.4 we perform two exercises. First, we analyze the effects of the MTO policy using a version of the model with only two neighborhoods and compare that to our baseline results. Second, we consider a different implementation of the MTO policy, where voucher recipient families are required to move to the middle-income neighborhood, in which case the richer families who used to live in the targeted neighborhood might decide to move to the richer one.

the general equilibrium sorting effects that arise in our model when neighborhood spillovers endogenously change in response to policies.

From a more micro perspective, two other related papers studying the MTO experiment using a location-choice model are Galiani et al. (2015) and Davis et al. (2021). Galiani et al. (2015) use the MTO experiment to identify structural parameters of a location-choice model and run a counterfactual to understand the role of mobility counseling in increasing take-up rates. They then use the quantitative model to study the effect of restricting the criteria to use the voucher. In particular, they show that forcing voucher receivers to move to neighborhoods with even lower poverty rates, would reduce the take-up rates and possibly backfire. In a similar spirit, Davis et al. (2021) estimate preferences in a dynamic location-choice model and emphasize the trade-off of restricting criteria to use the voucher to move to lower poverty neighborhoods: on the one hand, households who use the voucher end up in better locations, but on the other hand, fewer families end up taking up the voucher. They are closer to our paper in focusing on the general equilibrium effect of the policy. In particular, they emphasize a trade-off of restricting the use of the voucher to lower poverty neighborhoods: the endogenous increase in rental rate in high opportunity locations induces some households who used to live there to move to lower opportunity areas, dampening the net benefit effects of the policy.

There is a number of papers that use general equilibrium models similar to ours to study different types of neighborhood-specific policies. Among others, Rossi-Hansberg et al. (2010) propose a model of housing externalities, where the value of land is affected by the houses in the surrounding areas that affect the intensity of non-market interactions. They estimate how housing externalities decline with distance and use these estimates to evaluate the effect of the Neighborhoods-in-Bloom program, which is an urban revitalization program implemented in Richmond, Virginia, between 1999 and 2004.³ Aliprantis and Carroll (2018) propose a two-neighborhood model with residential choice and local human capital externalities. They then calibrate a version of the model with no residential choice to data from Chicago in 1960 to map the initial racial composition of the two neighborhoods, and use it to explore the effects of eliminating legal racial discrimination for neighborhood selection. In a similar spirit, Fogli et al. (2026c) explore the effects of removing covenant policies in a model with local spillovers and endogenous racial bias. Diamond et al. (2019) use a dynamic model with neighborhood choice disciplined with quasi-experimental variation in assignment of rent control to evaluate the benefits of rent controls for covered tenants and the welfare losses from decreased housing supply. More related to our paper, Agostinelli et al. (2020) focus on how parents affect the choice of peer groups. Their parenting style decision is close in spirit to our residential choice decision and is affected by the return to education. In a similar spirit to the MTO, they use this model to explore the effects of policy interventions to move children to better neighborhoods. Agostinelli et al. (2021) develop and estimate a spatial equilibrium model with residential sorting and school choice to study the effects of school bus transportation to have talented children from poor neighborhoods attending high-quality schools in richer neighborhoods. As in our paper, they have an endogenous school quality as well as heterogeneous skill distribution, and they show that scaling up the policy may have general equilibrium effects that dampen

³There is a vibrant empirical literature developing different strategies to estimate the strength of neighborhood externalities, including Redding and Sturm (2016) and Diamond and McQuade (2019) among others.

its beneficial effects. Gaubert et al. (2021) explore the efficiency and equity consequences of place-based redistribution policies.

Our paper is connected to a large empirical literature that studies the effects of the MTO program on both non-economic and economic objects. Initial research, including Katz et al. (2001), Kling et al. (2007), Clampet-Lundquist and Massey (2008), Ludwig et al. (2013), did not find a significant impact of the housing vouchers on the future income and employment status of the children of the voucher recipient families. More recently, Chetty et al. (2016), using administrative data including the adult income of individuals who were children when their families participated in the program, have found significant economic effects for children who moved to lower poverty neighborhoods when they were below 13 years old. Bergman et al. (2019) use a randomized controlled trial with housing voucher recipients in Seattle and King County to show that barriers in the housing search process are an important driver of residential segregation by income.

2 Model

In this section, we present a general equilibrium model with residential choice and local spillovers, building on Fogli et al. (2026b). Parents choose the neighborhood where to live, taking into account that different local spillovers affect their children's future income. A key element of the model is that the local spillovers are endogenous and depend on the sorting of families into neighborhoods. We also assume that a skill premium shock is the primary source of the rise in inequality.

2.1 Environment

The economy consists of three neighborhoods, denoted by $k \in \{A, B, C\}$, and is populated by overlapping generations of agents who live for two periods. An agent is a child in the first period and a parent in the second period. A parent at time t earns a wage $w_t \in [\underline{w}, \bar{w}]$, has one child with ability $a_t \in [\underline{a}, \bar{a}]$, and is born in neighborhood $b_t \in \{A, B, C\}$. The ability of a child is correlated with the ability of the parent. In particular, $\log(a_t)$ follows an AR1 process

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

where v_t is normally distributed with mean zero and variance σ_v , $\rho \in [0, 1]$ is the autocorrelation coefficient, and x is a constant normalized so that the mean of a_t is equal to 1. The joint distribution of parents' wages, children's abilities, and birth neighborhoods evolves endogenously and is denoted by $F_t(w_t, a_t, b_t)$, with $F_0(w_0, a_0, b_0)$ taken as given.

There is a continuum of landlords who build houses every period and rent them out to the parents. At the end of each period the houses fully depreciate. All houses are of the same dimension and quality and the rent in neighborhood k at time t is denoted by R_{kt} . The construction cost in neighborhood k at time t is equal to $\mathcal{C}(H_{kt}) = \phi_k / (1 + \phi_k) \lambda_k^{-1/\phi_k} H_{kt}^{(1+\phi_k)/\phi_k}$, where H_{kt} is the supply of houses in neighborhood k at time t .

The landlord's profits at time t are then equal to

$$\Pi_t = \sum_{k=A,B,C} \{R_{kt}H_{kt} - \mathcal{C}(H_{kt})\}.$$

The landlord's maximization problem generates a standard upward-sloping housing supply curve for each neighborhood k at time t , given by

$$H_{kt} = \lambda_k R_{kt}^{\phi_k},$$

where ϕ_k represents the housing elasticity in neighborhood k , and λ_k is a shift parameter in the same neighborhood.

At the end of each period, the landlords pool their profits Π_t and redistribute them to the parents with income in the top x -th percentile of the city-wide income distribution, in proportion to their income. Parents with wage $w_t \geq \bar{w}_t$, where $F_t(\bar{w}_t) = 1 - x$, get a fraction of the total profits Π_t equal to $s(w_t, a_t, b_t) = f_w(w_t)w_t / \int_{\bar{w}_t}^{\infty} f_w(w_t)w_t dw_t$. From now on we pick $x = .8$.⁴

A parent with income w_t , child's ability a_t , and birth neighborhood b_t chooses consumption, $c_t(w_t, a_t, b_t)$, the neighborhood where to raise their children, $n_t(w_t, a_t, b_t)$, and how much to invest in their children's education, $e_t(w_t, a_t, b_t)$. Parents who choose a neighborhood n_t different from their birth neighborhood b_t , incur in a utility cost μ . The educational choice is continuous and the cost of education is equal to τe_t^γ , with $\tau > 0$ and $\gamma > 0$.

The future wage of a child with parent (w_t, a_t, b_t) who grows up in neighborhood n_t and receives education e_t is equal to

$$w_{t+1} = \Omega(w_t, a_t, e_t, S_{n_t}, \varepsilon_t) \equiv (y + a_t e_t \eta_t (\beta_0 + \beta_1 S_{n_t})^\xi) w_t^\alpha \varepsilon_t, \quad (1)$$

where ε_t is an iid noise with cdf Ψ , normally distributed with mean one and standard deviation σ_ε , $\eta_t > 0$ captures a skill premium shock, and S_{n_t} is the strength of a local spillover in neighborhood n_t at time t that affects the returns to education. The wage equation shows that a child's wage is increasing in her ability, in her education level, in her parent's income, and is affected by the neighborhood where she grows up because of the local spillover.⁵

The strength of the spillover effect in neighborhood k at time t is equal to the expected future income of the children growing up in that neighborhood:

$$S_{kt} = \frac{\int \int_{n_t(w_t, a_t, b_t)=k} W_{t+1}(w_t, a_t, b_t, \varepsilon_t) dF_t(w_t, a_t, b_t) d\Psi_t(\varepsilon_t)}{\int_{n_t(w_t, a_t, b_t)=k} dF_t(w_t, a_t, b_t)}, \quad (2)$$

where we define $W_{t+1}(w_t, a_t, b_t, \varepsilon_t) \equiv \Omega(w_t, a_t, e_t(w_t, a_t, b_t), S_{n_t(w_t, a_t, b_t)t}, \varepsilon_t)$, that is, the children's wage given optimal education and residential choices of the parents. Given that wages are increasing in ability and in parents' wage, neighborhoods with higher spillovers are neighborhoods with both richer parents and children

⁴The results are similar for different values of x .

⁵Parents' wages affect children's wages both directly and indirectly through the educational and residential choices.

with higher ability. This formalization allows us to capture a number of different sources of local spillovers, from pecuniary ones, such as the quality of public schools, to non-pecuniary ones, such as social norms, peer effects, information externalities and so forth. We chose this general specification because we are going to quantitatively discipline the strength of the local externalities with the estimates of neighborhood exposure effects in Chetty and Hendren (2018), who do not distinguish among different channels. From now on, we define the neighborhood with the highest and lowest spillover as $\bar{n}_t \equiv \arg \max_{k \in A, B, C} S_{kt}$ and $\underline{n}_t \equiv \arg \min_{k \in A, B, C} S_{kt}$.

Parents care about their own consumption and about their child's future income.⁶ We also introduce two types of preference shocks over neighborhoods to capture the role of fixed amenities (e.g., waterfront, parks, and so forth) as well as idiosyncratic preference for different locations (e.g., family network). We assume that neighborhood A has better amenities than B, and B has better amenities than C and that only a random fraction π of parents care about amenities. Hence, the utility from consumption for a parent who chooses neighborhood n is given by $\log(\theta_n c)$, where $\theta_A > \theta_B = 1 > \theta_C$ with probability π and $\theta_A = \theta_B = \theta_C = 1$ with probability $1 - \pi$.

The preferences of parent (w_t, a_t, b_t) can be written as

$$\log(\theta_{n_t} c_t) + \log(w_{t+1}) + \sigma_\zeta \zeta_{n_t} - \mu I_{n_t \neq b_t}$$

where ζ_{n_t} is the idiosyncratic shock, which follows a Type-I Extreme Value distribution with scale parameter σ_ζ . This shock introduces some additional randomness that is not systematically related to some particular ranking of the neighborhoods and helps making the model analytically tractable. Moreover μ is the utility cost of moving neighborhood.

We assume that the government runs a public assistance (PA) program for the parents with low income to ensure everybody a minimum level of consumption \underline{c} . The program is financed with a proportional income taxation at the city level with tax rate κ_t . Parents are eligible if they have income below the cut-off $\hat{w}_t = (R_{\underline{n}_t} + \underline{c}) / (1 - \kappa_t)$. Eligible parents have to live in the neighborhood with the lowest spillover \underline{n}_t . If their income is higher than the rental rate in that neighborhood, they have to pay for it and they will receive a consumption transfer (e.g. food stamps) such that they achieve a minimum level of consumption \underline{c} . If their income is below the rental rate, the government will cover the remaining portion of the rent and transfer \underline{c} to them. To sum up, a parent with $w_t < \hat{w}_t$ will receive a transfer in terms of consumption equal to $T_t(w_t) = \underline{c} - \max\{(1 - \kappa_t)w_t - R_{\underline{n}_t}, 0\}$.⁷ For each eligible parent with $w_t < \hat{w}_t$, the government will need resources equal to $R_{\underline{n}_t} + \underline{c} - (1 - \kappa_t)w_t$.⁸ The government's budget balance condition at each time t is then given by

⁶This simplifying assumption is common in this class of models. The assumption that agents cannot borrow is for realism, given that typically people cannot borrow against children's future income.

⁷Given that this transfer is in terms of consumption, a parent with transfer $T_t(w_t)$ will have to choose $c(w_t, a_t, b_t) \geq T_t(w_t)$.

⁸If the eligible parent has a wage smaller than the rental rate, then the government has to pay for the residual rental rate on top of the transfer.

$$\int_{w_t < \hat{w}_t} (R_{n_t} + \underline{c} - (1 - \kappa_t)w_t) dF_t(w_t, a_t, b_t) \leq \kappa_t \left(\int w_t dF_t(w_t, a_t, b_t) \right). \quad (3)$$

To summarize, the optimization problem for a parent (w_t, a_t, b_t) with $w_t \geq \hat{w}_t$ at time t can be written as:

$$\begin{aligned} U(w_t, a_t, b_t) &= \max_{c_t, e_t, n_t} \log(\theta_{n_t} c_t) + E_t[\log(w_{t+1})] + \sigma_\zeta \zeta_{n_t} - \mu I_{n_t \neq b_t} \\ \text{s.t. } c_t &\leq (1 - \kappa_t)w_t + s_t(w_t) \Pi_t I_{w_t > \bar{w}_t} - R_{n_t} - \tau e_t^\gamma \\ w_{t+1} &= \Omega(w_t, a_t, e_t, S_{n_t}, \varepsilon_t), \end{aligned} \quad (\text{P1})$$

taking as given spillovers and rental rates, S_{n_t} and R_{n_t} for $n_t = A, B, C$.

Instead a parent (w_t, a_t, b_t) with $w_t < \hat{w}_t$ will have to live in the neighborhood with the lowest spillover $n_t(w_t, a_t, b_t) = \underline{n}_t$, and her optimization problem will be⁹

$$\begin{aligned} U(w_t, a_t, b_t) &= \max_{c_t, e_t} \log(\theta_{\underline{n}_t} c_t) + E_t[\log(w_{t+1})] + \sigma_\zeta \zeta_{\underline{n}_t} - \mu I_{\underline{n}_t \neq b_t} \\ \text{s.t. } c_t &\leq (1 - \kappa_t)w_t - \min\{(1 - \kappa_t)w_t, R_{\underline{n}_t}\} - \tau e_t^\gamma + T_t(w_t) \\ T_t(w_t) &= \underline{c} - \max\{(1 - \kappa_t)w_t - R_{\underline{n}_t}, 0\} \\ c(w_t, a_t, b_t) &\geq T_t(w_t) \\ w_{t+1} &= \Omega(w_t, a_t, e_t, S_{\underline{n}_t}, \varepsilon_t). \end{aligned} \quad (\text{P2})$$

All parents with after-tax income below the rental rate in the neighborhood will have consumption $c_t(w_t, a_t, b_t) = \underline{c}$, and no resources for education, so $e_t(w_t, a_t, b_t) = 0$. Parents who have leftover income after paying the rent could choose some positive level of education.

In our analysis, we assume that there are complementarities between the spillover and ability, education, and wages, and between education and ability, and education and wages.

Assumption 1 *The composite function $\Omega(w, a, e, S, \varepsilon)$ has increasing differences in a and S , in e and S , in w and S , in a and e , and in w and e .*

These complementarities assumptions drive residential segregation by income and ability in equilibrium. Although it is hard to get direct estimates of innate ability, the complementarity between innate ability and education and between innate ability and neighborhood spillover reflect some of the findings of the recent empirical literature.¹⁰

⁹Notice that there are no landlords' profits in the budget constraint because $\bar{w}_t < \hat{w}_t$.

¹⁰Our assumptions of complementarity between innate ability and education and between innate ability and neighborhood spillover are consistent with the latest research on technology of skill formation. Cunha et al. (2010) show that the higher the initial conditions for cognitive and non-cognitive skills of children, the higher the return to parental

2.2 Equilibrium

We are now ready to define an equilibrium.

Definition 1 For a given initial wage distribution $F_0(w_0, a_0, b_0)$, an equilibrium is characterized by a sequence of educational and residential choices, $\{e_t(w_t, a_t, b_t)\}_{t \geq 0}$ and $\{n_t(w_t, a_t, b_t)\}_{t \geq 0}$, a sequence of rents and spillover's sizes in the three neighborhoods, $\{R_{kt}\}_{t \geq 0}$ and $\{S_{kt}\}_{t \geq 0}$ for $k = A, B, C$, a sequence of tax rates $\{\kappa_t\}_{t \geq 0}$, and a sequence of distributions $\{F_t(w_t, a_t, b_t)\}_{t \geq 0}$ that satisfy:

1. agents' optimization: for each t and given R_{kt} and S_{kt} for $k = A, B, C$
 - (a) for (w_t, a_t, b_t) with $w_t \geq \hat{w}_t$, e_t and n_t solve problem (P1);
 - (b) for (w_t, a_t, b_t) with $w_t < \hat{w}_t$, $n_t = \underline{n}_t$ and e_t solves problem (P2);
2. spillovers' consistency: for each t , equation (2) is satisfied for $n = A, B, C$;
3. market clearing: for each t and $k \in \{A, B, C\}$, R_{kt} is such that

$$\lambda_k \left(\frac{R_{kt}}{\bar{w}_t} \right)^{\phi_k} = \int_{n_t(w_t, a_t, b_t)=k} dG_t(w_t, a_t, b_t); \quad (4)$$

4. wage dynamics are consistent with optimal choices: for each t ,

$$w_{t+1} = \Omega(w_t, a_t, e_t(w_t, a_t, b_t), S_{n_t(w_t, a_t, b_t)t}, \varepsilon_t). \quad (5)$$

5. budget balance: for each t , κ_t is such that equation (3) is satisfied.

In equilibrium, the presence of local spillovers, together with assumption 1, generate residential segregation by income and ability.¹¹ Talented children from poor families who grow up in the poorer neighborhoods do not have the same opportunities as children of richer families who can afford a neighborhood with higher spillover. Moreover, since higher local spillovers increase the return to education, parents who live in poorer neighborhoods, everything else equal, end up investing less in education than parents who live in richer neighborhoods. This amplifies future inequality and segregation and further reduces intergenerational mobility.¹²

investment in children at later stages in life. As they highlight, "Family environments and genetic factors may influence these initial conditions." In our model, parental investment in children's future outcomes takes place both through educational investment and through residential choice. Moreover, the recent human capital literature, reviewed in Sacerdote (2011), also highlights the presence of non-linearity in peer effects, which are one of the forces behind our spillover effects. In particular, Sacerdote (2001), Imberman et al. (2012), and Lavy et al. (2012) find that high ability students are the ones who benefit the most from peer effects of other high ability students. Another paper that speaks more specifically to the complementarity between ability and spillover effects is Card and Giuliano (2016), which shows that high achievers from minority and disadvantaged groups show high returns when included in school tracking programs.

¹¹Residential segregation by income is also driven by the presence of local amenities.

¹²If the complementarity is strong enough, the model could potentially generate multiple equilibria. However, for

3 Calibration

In this section, we describe the calibration strategy we use for our quantitative analysis.

First, we assume that our economy represents a representative US metro area. We also assume that the economy is in steady state in 1980. Our baseline model features the public assistance program, but does not include any neighborhood-specific policy. We introduce a skill premium shock in 1990 so that the baseline model can generate an increase in inequality after 1980, close to the data. In particular, we choose the size of this shock in order to match the increase in the college premium between 1980 and 1990. In our exercises, we will compare the dynamic response to this shock of versions of the economy with different neighborhood-specific policies.

Table 1 shows the targets of our baseline calibration, which we are now going to discuss.

In order to map our model to the data, we interpret one period as ten years.¹³ Moreover, to define the three neighborhoods, we rank census tracts in each MSA according to the number of families living there who belong to the top 20th percentile of the income distribution (“rich families” from now on). For each MSA, we define neighborhood A as the set of census tracts with more than 30% rich families, neighborhood C as the set of tracts with less than 17% rich, and neighborhood B as the residual. Then, we average the MSA-specific neighborhoods A, B, C across all MSAs in our sample to construct the representative neighborhoods A, B, and C in our model. The cut-offs chosen to define the neighborhoods imply that in 1980 roughly 50% of the US population lives in neighborhood C and the rest is roughly split between A and B.¹⁴

The first three targets in Table 1 are measures of residential segregation by income and income inequality at the metro area level. For all these measures, we restrict the sample to families with children because our mechanism emphasizes the parental decision to invest in the children’s education. First, we target the 1980 value of the dissimilarity index by income as a measure of segregation. We calculate the dissimilarity index for each MSA and then take the average, weighting by population.¹⁵ Second, we target the value of the Gini in 1980 as a measure of income inequality. Using Census data, we calculate the Gini coefficient for each metro area and then we average them, weighting by population. Third, as an additional measure of income inequality, we target the ratio of the average income for families in the top 25th percentile of the income distribution to the average income for families in the bottom 25th percentile.

In addition, we target the level of intergenerational mobility, measured as the rank-rank correlation between

the parameter values we obtain in our calibration, we always numerically converge to the same steady-state equilibrium.

¹³Following Fernandez and Rogerson (1998), we make this choice motivated by our focus on parental education investment, as school duration extends to 10 or 15 years, depending on which level of education one targets. Another factor in our choice of 10 years is that census data are available every 10 years.

¹⁴This definition of C allows us to have enough room to expand the policy to a progressively larger number of families targeted by the program.

¹⁵We define the dissimilarity using rich and poor as the mutually exclusive groups, where rich are the families in the top 20th percentile of the MSA income distribution and poor are all the others. Moreover, we use the MSA-specific neighborhoods A, B, and C, constructed as described above, as the geographic subunit of analysis.

Table 1: *Calibration Targets*

Description	Data	Model	Source
Dissimilarity index by income	0.334	0.333	Census 1980
Gini coefficient	0.376	0.377	Census 1980
Income 25th/75th p	0.667	0.712	Chetty and Hendren (2018)
Rank-rank correlation	0.335	0.352	Chetty et al. (2014b)
Return to spillover 25th p	0.06	0.06	Chetty and Hendren (2018)
Return to spillover 75th p	0.05	0.05	Chetty and Hendren (2018)
Return to college 1980	0.391	0.397	Goldin and Katz (2009)
Return to college 1990	0.549	0.563	Goldin and Katz (2009)
Neighborhood A size 1980	0.194	0.193	Census 1980
Neighborhood A size 1990	0.217	0.209	Census 1990
Neighborhood B size 1980	0.301	0.301	Census 1980
Neighborhood B size 1990	0.250	0.277	Census 1990
Share of rich in A 1980	0.437	0.444	Census 1980
Share of rich in B 1980	0.225	0.227	Census 1980
College share A 1980	0.340	0.336	Census 1980
College share B 1980	0.178	0.211	Census 1980
Rent ratio R_A/R_B 1980	1.253	1.257	Census 1980
Rent ratio R_B/R_C 1980	1.277	1.279	Census 1980
Average rent ratio 1990	1.300	1.304	Census 1990
Average city housing elasticity	1.75	1.75	Saiz (2010)
SNAP population share 1980	0.07	0.08	USDA Food Nutrition Service

log wages of parents and children estimated using administrative records by Chetty et al. (2014b). In particular, they use children born between 1980 and 1982, calculate parental income as mean family income between 1996 and 2000 and children’s income as mean family income between 2011 and 2012, when the children are approximately 30 years old. Given that this correlation is calculated over several decades, we map it in the model to the average rank-rank correlation across 1980, 1990, and 2000.

An important target for our exercise is the “return to spillover”, that is, the effect of the neighborhood exposure on children’s income in adulthood. To measure this effect, following Fogli et al. (2022) we rely on the results from the quasi-experiment in Chetty and Hendren (2018), who use tax returns data for all children born between 1980 and 1986, to estimate the effect of local spillovers on children’s future income, by looking at movers across US counties.¹⁶ We target their estimates which imply that for a child with parents at the 25th (75th) percentile of the national income distribution, growing up in a 1 standard deviation better county from birth would increase household income in adulthood by approximately 6.2% (4.6%).¹⁷ In our model, we consider parents at the 25th percentile and at the 75th percentile of the income distribution

¹⁶Chetty and Hendren (2018) control for selection effects by looking at families who move from one county to another with kids of different age, so that they were exposed for different fractions of their childhood to the new county. We focus on their estimations for families moving across counties within the same commuting zone, given that we use the metro area as our geographic unit of analysis.

¹⁷See table II in Chetty and Hendren (2018).

who decide to live in a neighborhood different from their birth neighborhood and calculate the ratio of the standard deviation of the expected future wage of their children divided by the average wage of the parents. Given the timing of the estimates from the data, we again map these numbers to the average “spillover effects” in the model across 1980, 1990, and 2000.

As we mention above, we target the increase in US college premium between 1980 and 1990 (from Goldin and Katz (2009)). In the model, we map the skill premium to the steady state difference between the average log wage of college-educated individuals and the average log wage of the others. To define college-educated individuals, given that the educational choice is continuous, we define a cut-off \hat{e} 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 Census data and keep it constant.¹⁸

In our model, the sizes of the three neighborhoods are endogenous. We use micro data on census tracts’ population, to target the size of the three neighborhoods both in 1980 and 1990. It is useful to use both years, to recover both the housing supply shifters and the housing supply elasticities for the three neighborhoods. We also target the share of rich families in the different neighborhoods in 1980, where rich are again defined as families in the top 20th percentile of the income distribution. Moreover, we use Census tract data to calculate the share of people over 25 years old who completed college residing in neighborhood A, B, and C for the average MSA.

Two other key objects in the model are the ratio of rental rates in neighborhood A to neighborhood B and in neighborhood B to neighborhood C in 1980. We use housing values in 1980 at the census tract level from the Census data and convert them into rental rates.¹⁹ We also target the average ratio of rental rates (A/B and B/C) in 1990 so that the growth rate in rental rates help us pinning down the moving costs. Moreover, we target the average MSA housing elasticity, using the estimate in Saiz (2010).²⁰

Finally, we map the minimum level of consumption guaranteed by our public assistance policy to the food stamps in the Supplemental Nutrition Assistance Program (SNAP). In 1980, about 7% of US households were receiving SNAP benefits.²¹

Table 2 shows our calibrated parameters, their calibrated value, and their description. We normalized the mean of the ability process a_t , the mean of the noise shock of the wage process ε_t , the education cost τ , the steady state college return parameter η , and the steady state rental rate in neighborhood C.²²

¹⁸To calculate this number, we look at the number of people over 25 year old who completed college at the census tract level.

¹⁹We use a standard coefficient of 0.05 for the conversion.

²⁰To validate our approach, we also draw on measures of within-city heterogeneity in elasticities across neighborhoods, as reported by Baum-Snow and Han (2024). Specifically, they highlight significant variation in housing supply elasticity based on proximity to central business districts (CBDs). Using their estimates, we compute the mean-to-standard deviation ratio of housing supply elasticity for neighborhoods located at the CBD, halfway to the region’s edge, and at the region’s edge, which yields a value of 1.08. Similarly, we calculate the mean-to-standard deviation ratio for our calibrated elasticities, which also results in a value of 1.08.

²¹See Pew Research Center analysis of data from USDA Food and Nutrition Service.

²²For details about the normalizations, see Appendix C.1.

Table 2: *Parameters*

Parameter	Value	Description
ρ	0.47	Autocorrelation of log ability
σ_v	1.11	St. dev. of log ability
σ_ε	0.60	St. dev. of log wage noise shock
α	0.24	Wage function parameter
β_0	0.15	Wage function parameter
β_1	0.15	Wage function parameter
ξ	1.16	Wage function parameter
y	1.04	Wage fixed component for no-college
γ	5.70	Education cost parameter
θ_A	1.28	Preference shock for neighborhood A
θ_C	0.42	Preference shock for neighborhood C
π	0.46	Preference shock probability
σ_ζ	0.20	St. dev. of idiosyncratic preference shock
λ_A	0.68	Shift parameter of housing supply in A
λ_B	0.31	Shift parameter of housing supply in B
ϕ_A	1.94	Elasticity of housing supply in A
ϕ_B	0.05	Elasticity of housing supply in B
ϕ_C	3.26	Elasticity of housing supply in C
η	2.81	Skill premium 1990
μ	0.29	Moving cost
\underline{c}	0.004	Minimum Consumption

Moreover, we feed the empirical time series for the population growth of the average metro area. In particular, we set the population growth in each decade from 1980 to 2020 to be equal to the empirical counterpart from the CBO data and to the projections from the CBO from 2030 onwards.²³

Our calibrated model is able to generate dynamic patterns for income inequality and segregation by income roughly in line with the data. In particular, it matches very closely the gini index, while it generates slightly higher growth in the dissimilarity index. The gini grows by 25% between 1980 to 2010 in the model relative to 24% in the data, while the dissimilarity index grows by 40% in the model and by 33% in the data. Recall that by calibration design, we match the level of inequality and segregation in 1980 and the increase in the college premium between 1980 and 1990, which we consider the primary source for the increase in inequality in the data. Our model generates an amplification effect due to the endogenous feedback between inequality and segregation. As the skill premium increases, the return to live in neighborhoods with higher spillover increases, pushing up the house prices in those neighborhoods and generating more residential segregation by income. In turn, this further amplifies future inequality.

²³We adjust the population growth values from the CBO that cover the whole US population to match the average population growth between 1980 and 2010 in our sample of metros using Census data. For details, see Appendix C.

4 Moving to Opportunity and Place-Based Transfer Programs

We now use our calibrated model to explore the aggregate effects of two alternative neighborhood-specific policies: the moving-to-opportunity (MTO) and the place-based transfer (PBT) programs. In particular, we introduce each policy as an unexpected permanent shock in 1990. We start from exploring the impact response of the economy to these policies and move to the analysis of their dynamic effects in Section 6. In the next section, we introduce another type of neighborhood specific policy: the place-based investment (PBI) program.

4.1 MTO Program

The MTO program was a randomized housing-mobility experiment run by the US Department of Housing and Urban Development (HUD) in Baltimore, Boston, Chicago, Los Angeles, and New York in 1994–1998. Eligible families were those with children under 18 living in high-poverty neighborhoods (i.e., with poverty rate above 40% in 1990) and in public or project-based assisted housing for low-income families (i.e., with income below 50% of the local median). After preliminary screening, eligible families that decided to apply to the program were randomly assigned to one of three groups: (i) an *experimental* group, receiving a housing voucher that could only be used to move to low-poverty areas (1990 poverty below 10%); (ii) a standard Section 8 group, receiving a housing voucher with no location constraint; or (iii) a control group, receiving no assistance through the MTO program, but which continued to be eligible for housing assistance and other welfare programs. Voucher holders paid 30% of household income toward rent and utilities, with the remainder covered by the voucher up to Fair Market Rent (defined as the 40th percentile of rental costs in a metro area).²⁴

Chetty et al. (2016) documented that this program succeeded in improving the long-term outcomes of children younger than 13 years old.²⁵ Given this finding, a natural question is whether this program can be scaled up to further reduce inequality and improve intergenerational mobility. While for the small scale experiment, the general equilibrium effects are negligible, the natural challenge when we scale up the policy is to consider that the general equilibrium effects may become sizeable and affect the efficacy of the program.

We now introduce a policy resembling the MTO program in 1990 as a permanent unexpected shock, together with the skill premium shock. To map the MTO program to our model, we impose that the families enrolled in the program at time t need to satisfy two criteria: 1) they need to live in the neighborhood with the lowest spillover, and 2) they need to have an income below a given threshold \hat{w}_t .²⁶ In our main exercises, we focus on the experimental group and assume that the enrolled families who decide to accept the housing voucher at time t need to move to the neighborhood, which had the highest spillover in the previous period. Moreover,

²⁴For more details about the program, see Appendix A.

²⁵Early evaluations reported no substantial effects on children’s later earnings or employment, though improvements were found for health and well-being (Katz et al., 2001; Kling et al., 2007; Clampet-Lundquist and Massey, 2008; Ludwig et al., 2013).

²⁶See Appendix B for more details on the introduction of the policy in the model.

if they do, they have to contribute a share $1 - q$ of their income to pay for rents and the government pays the difference up to a voucher cap \bar{r}_i . In our calibrated steady state in 1980, neighborhood A is the one with the highest spillover and neighborhood C the one with the lowest. We also consider an alternative way of mapping the experimental group case to the model by assuming that the recipient families need to move to the neighborhood with the intermediate spillover, neighborhood B in steady state, instead of A. We show some of these results in this section and a more complete analysis of this case in Appendix D. In the same appendix, we also explore the effects of section 8, that is, the policy where recipient families can choose which neighborhood to move to.

Let us first describe how we set the enrollment percentile p , such that $F(\hat{w}_{1990}) = p$. First, we calculate the average share of enrolled families in the five cities where the MTO was introduced, by dividing the number of individuals effectively enrolled in the MTO program by the total population. The total population in the five cities was 13,902,026 and there were 4,142 families enrolled in the program.²⁷ We assume that the average family size is 4 and obtain a share of enrolled individuals equal to 0.001. Given that the eligibility criteria also require to live in the neighborhood with the lowest spillover, the enrollment share 0.001 translates into an enrollment percentile out of the city-wide income distribution p equal to 0.15%. Table 3 shows the mapping between enrollment shares out of the total population and enrollment percentiles out of the city-wide income distribution over time as implied by the model dynamics. As the table shows, the fraction of families satisfying the enrollment criterion for a given p varies endogenously over time because of the endogenous change of the income distribution and the endogenous residential choices.²⁸ From now on, we will refer to the scale of the policy as the enrollment percentile p .

Table 3: MTO enrollment shares for different p

	Enrollment Shares					
	p					
	0.15%	3%	6%	9%	12%	15%
1990	0.001	0.019	0.039	0.053	0.071	0.105
2000	0.001	0.021	0.040	0.053	0.068	0.097
2010	0.001	0.020	0.038	0.059	0.076	0.106

Moreover, we set q and \bar{r} , to match the MTO program. The program imposes families receiving the housing voucher to pay 30% of their income for rents, so we set $q = 70\%$. The program prescribes a housing voucher cap between 80% and 100% of the local Fair Market Rent (FMR) in that period, which is defined as the 40th percentile of rents for all rental units of a given bedroom size in a given MSA. We take the extreme of 100% of the FMR and calculate the value of the 40th percentile of rents in the data relative to the rent in A in 1990

²⁷See Table A1 in Appendix A.

²⁸Children of families enrolled in the program at time t are not necessarily eligible for the program at time $t + 1$.

based on our neighborhood definition. This is roughly equal to 80% of the rental rate in neighborhood A in 1990. We then set $\bar{r}_t = .8 * R_{At}$.

Panel (a) in Table 4 summarizes the value and description of the three parameters just described.

Table 4: *MTO Policy*

— Panel A. Parameters —			
Parameter	Value	Description	
p	15%	MTO enrollment income percentile	
q	0.70	MTO income share after down-payment	
\bar{r}	0.80	Voucher cap as a fraction of R_A	
— Panel B. Validation —			
Description	Data	Model	Source
Voucher take-up rate	0.47	0.48	MTO manual
Income gain of recipient children	0.16	0.17	Chetty and Hendren (2018)

The model is able to generate a voucher take-up rate and an income gain for the children of voucher recipients at the time of the policy introduction that are roughly in line with the data. Of the 4,142 families that applied, only 47% ended up taking up the voucher, which is close to a take-up rate of 48% generated by the model.²⁹ Moreover, Chetty et al. (2016) estimate an increase in adult annual income of 31% for children younger than 13 at the time when their family took up the MTO voucher in the experimental group relative to the control group, while they find negligible or slightly negative effects for children older than 13. In our model, the income gain of the children of voucher recipient families is calculated as the percentage difference between the adult expected income of children in families who received the voucher relative to the adult expected income of the same children if there were no policy in place. Given that our model does not distinguish between young and old kids, we compare our model implied statistic with a simple average income gain of young and old kids in the data, which is approximately equal to 16%. Our model generates an income gain of 17%, which is a good validation of the model. Panel (b) in Table 4 summarizes the two statistics we use to validate the model.

4.2 Scaling up the MTO

We now use our model to study how the impact response of the economy to the introduction of the MTO program, that is, the response in 1990, varies if we scale up the program. To do so, we relax the enrollment requirement for voucher assignment by increasing the enrollment percentile p . We denote by $p_0 = 0.15\%$ the “effective scale” that corresponds to the effective MTO program and we scale it up to $p = 15\%$. It is important to notice that as we scale up the program, we also change the distribution of ability and income of the families who receive the voucher. In this section, we focus on the impact response (that is, the response

²⁹See Table A1 in Appendix A.

at the time of the policy introduction) of the economy to the policy for different scales. In Section 6, we will explore its dynamic effects.³⁰

As we explore the effects of scaling up the MTO program, it is important to consider the role of the endogenous change in the neighborhoods' spillovers and rental rates, due to the re-sorting of families across neighborhoods in response to the policy. While these general equilibrium effects are negligible for the effective scale of the program, they become larger the larger is the scale. Given that we describe the impact responses at time 1990, we can refer to the neighborhood with the highest and lowest spillovers as neighborhoods A and B respectively. Panel (a) in Figure 1 shows the impact response of the spillover to the introduction of the policy in the three neighborhoods, as a function of its scale. From now on, the pink star denotes the effective scale of 0.015. As the scale increases, there is a larger share of families who take up the voucher and move from neighborhood C to neighborhood A. Given that the voucher recipient families are poorer than the average families living in A and, given the correlation between income and ability, the spillover in A declines. At the same time, this improves the selection of families in neighborhood C and increases the spillover in C, reducing the gain of moving from C to A. As the spillover in neighborhood A decreases, more families decide to move from neighborhood A to neighborhood B, increasing the spillover in B. As we scale up the policy, such a process gets more pronounced until it can lead to a switch in the ranking of the neighborhoods' spillovers. In particular, when the cutoff is larger than the 12th percentile, the spillover in B becomes larger than the spillover in A, making B the most desirable neighborhood.³¹ Panel (b) in the same figure shows the average ability of families living in the three neighborhoods and confirms that the selection on ability is an important driver of the spillover response.

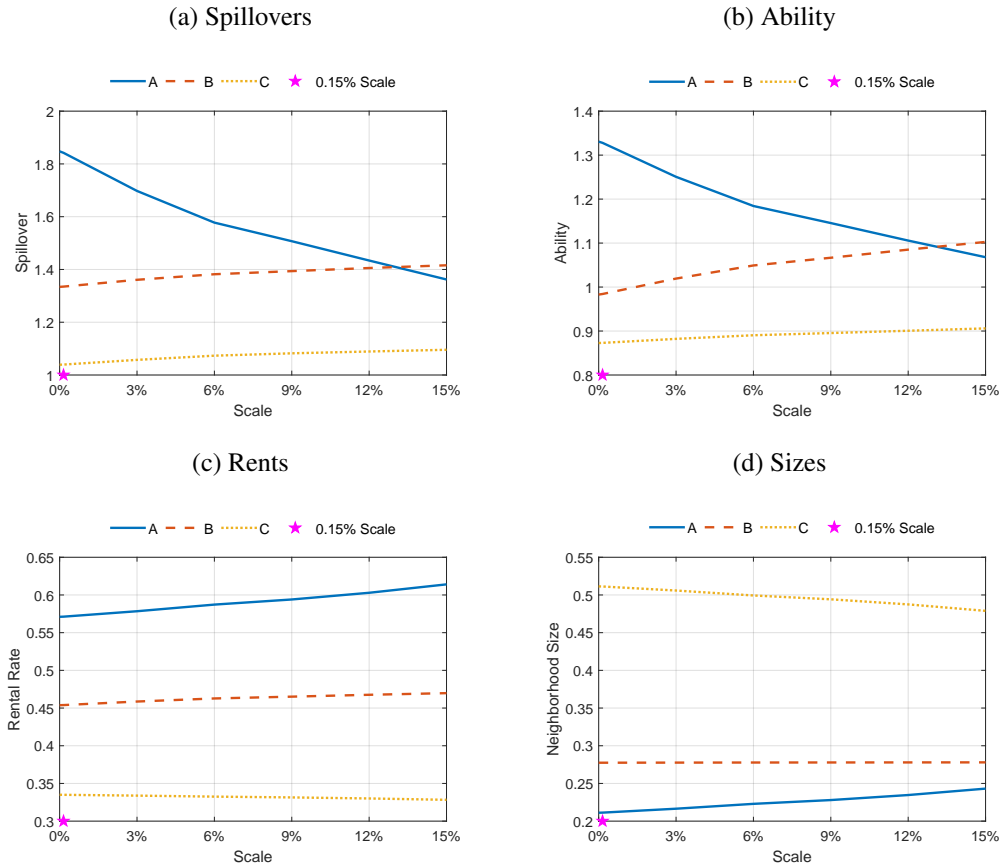
Moreover, panels (c) and (d) in the same figure show, respectively, the rental rates and the sizes of the three neighborhoods as a function of the policy's scale. The decline in the spillover of neighborhood A tends to reduce the demand to live there. However, the increase in demand because of the voucher recipient families moving to A dominates and makes the overall demand for neighborhood A higher. This translates into both an increase in the rental rate in A and an increase in the size of the neighborhood. The demand to live in B also increases because of the increase of families who decide to move from A to B, due to the fact that the gap in spillover declines, while the rental rate in A increases. This increases the rental rate and also slightly the size of neighborhood B. The demand to live in neighborhood C decreases because of the voucher receivers who are moving to A, although this effect is partially offset by the increase in the spillover due to the improvement of the composition of families living there. As a result, both the size and the rental rate in neighborhood C decline. The increase in rental rate in A and the decrease in rental rate in C further reduce the advantage of moving from C to A.³²

³⁰An alternative way of scaling up the policy is to assume that the eligible families are an increasing random proportion of families living in the neighborhood with the lowest spillover. This would eliminate selection effects, but would target a less poor portion of the population, while Section 8 eligibility criteria that need to be satisfied for the MTO seem to target the poorest families in the eligible pool. In Appendix D, we consider this alternative exercise.

³¹If the scale of the program increases further, the spillover in A further decreases to the point that for some parameters and a scale large enough, it can become even lower than the spillover in C.

³²Notice that the rental rate in C does not move much as C is the neighborhood with highest housing supply elasticity.

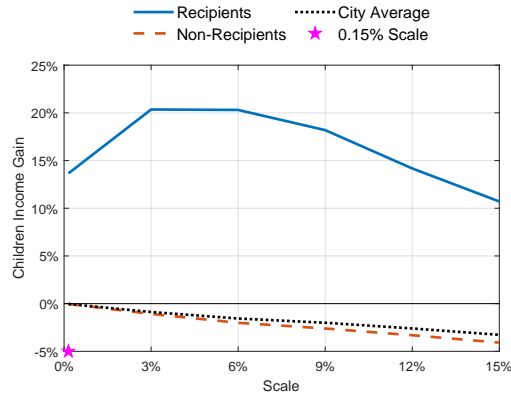
Figure 1: MTO - General Equilibrium Effects



To evaluate how the effectiveness of the policy changes with the scale, we first explore the policy’s impact on children’s future income. In particular, as the scale increases and the general equilibrium effects become sizeable, the policy is going to have effects on the future income not only of the children of recipient families, but also on all other children. The blue solid line in Figure 2 shows the income gain of the children of voucher recipient families, who benefit from growing up in a neighborhood with a higher spillover, as a function of the policy scale. As the scale increases, the average income gain increases at first, but then it starts to decline. This is due to different forces working in opposite directions. As the income cutoff of the enrolled families increases, richer families will be able to take up the voucher. On the one hand, given the positive correlation between wage and ability, this also means that recipient families will tend to have more talented kids. Due to the complementarity between local spillover and education, between spillover and parents’ wage, and between spillover and ability, richer families will have a larger income gain from moving to a neighborhood with higher spillover. Moreover, the presence of a cap on the voucher policy contributes largely to the increase in income gain going from p_0 to $p = 3\%$.³³ On the other hand, children

³³When the scale is small, a large fraction of the eligible families are so poor that they end up not taking up the

Figure 2: MTO - Effect on Children Income



of richer families have a smaller gain from the voucher policy because their parents would have invested in education even in the absence of the voucher. This second effect dominates when the scale of the policy becomes large enough, generating the inverted U shape. On top of these effects, as we scale up the policy, the general equilibrium effect that we have shown in Figure 1 further reduces the income gain of the voucher receivers' children, as the spillover gain from moving to A declines and the rental cost increases.

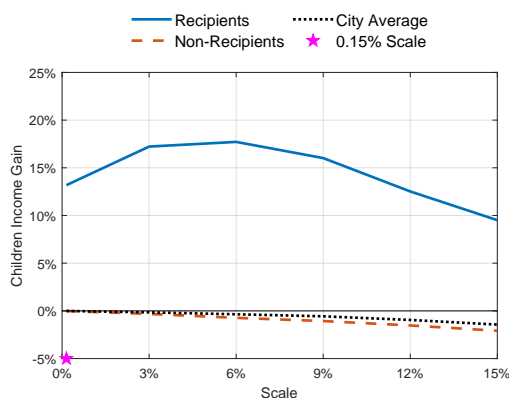
The red dashed line in Figure 2 shows how the policy affects the future income of children of families who are not voucher recipients. When the scale of the policy is equal to the effective scale, there is virtually no effect on the future income of children of non-recipient families, as the general equilibrium effects are negligible. However, as we scale up the policy and the general equilibrium effects kick in, children of non-recipient families suffer an income loss, the larger the larger is the scale. This is mainly due to children of families growing up in A, who are now exposed to a smaller spillover relative to when there is no policy. Moreover, these families have to pay higher rental rates and taxes, which reduce their investment in their children's education.

Finally, the black dotted line in Figure 2 shows the average children's income gain at the city level. For the effective scale of the program, the city-wide income gain is quite small but negative, because, although the recipient families enjoy a large income gain, they are a small fraction of the population and everybody else is paying taxes to cover the program, reducing investment in education. As the scale becomes bigger the city-wide average income loss becomes larger because of the general equilibrium effect just described.

In our baseline MTO exercise, recipient families are required to move to neighborhood A, which has the voucher. The reason is that the housing cost is higher than 30% of their income and is beyond the cap imposed by the policy. This means that these families would have to pay additional out-of-pocket expenses to move to A and a large fraction of them end up deciding not to take up the voucher, even if they have high ability children. When the scale increases, a larger fraction of families take up the voucher, so more high ability children have a chance to move to A and, given the complementarity between ability and spillover, the average income gain becomes larger.

highest spillover. A possible alternative interpretation of the “low poverty neighborhood” targeted by the actual MTO program is that families were instead required of moving to B. Figure 3 explores this alternative specification. The income gain for recipient children is lower than in the baseline case, peaking at around 18% compared to 22% when recipients move to A, reflecting the fact that the spillover in B is lower than in A. At the same time, the negative effects on non-recipient families are substantially milder, reaching only about -2% at the 15% scale compared to -5% when recipients move to A. This is because recipients are more similar to the families in B than to the families in A, and this is reflected in a smaller decline in the spillover in B. As a result, the city-wide average income loss is also smaller.

Figure 3: MTO - Effect on Children Income - Recipients move to B



An interesting implication of this specification is that now the richer families living in B might decide to move to neighborhood A, as the MTO voucher recipients move in B, driving down the local spillover. This tends to dampen the policy effect of reducing residential segregation by income. Figure D1 in Appendix D shows that residential segregation declines less when B is the targeted neighborhood relative to our benchmark.

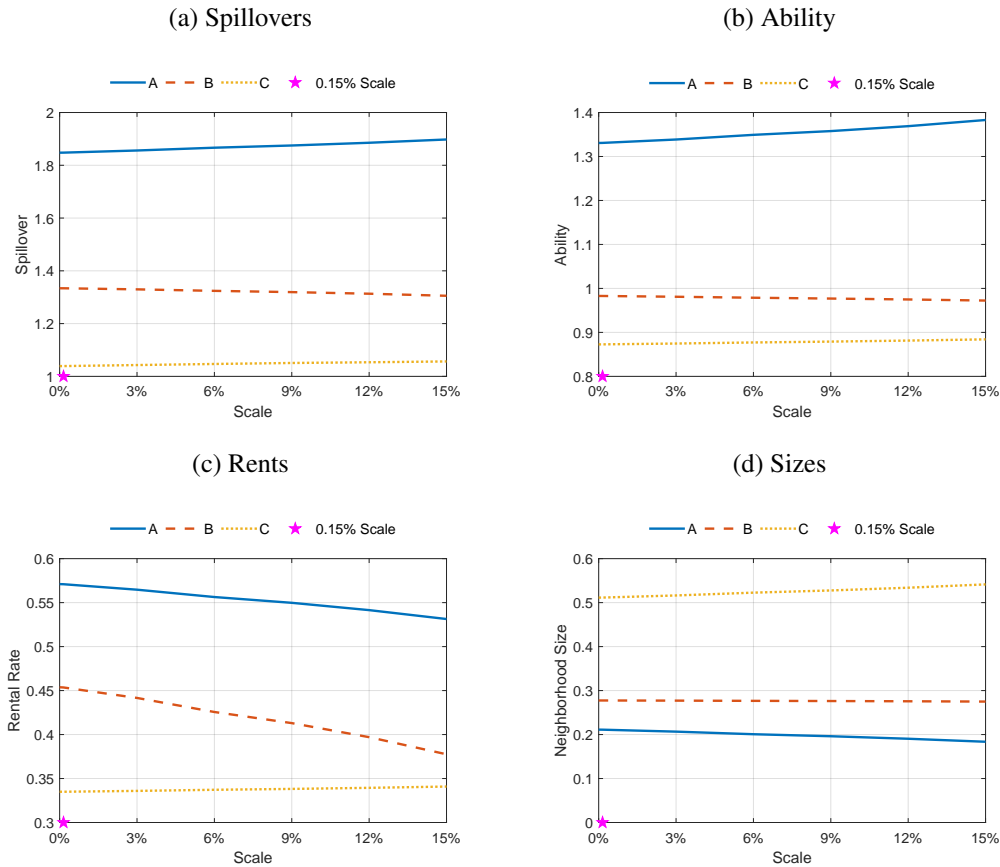
In Appendix D, we also provide additional results for this alternative specification, and we consider two other variants of the MTO policy: the case where recipients can choose whether to move to A or B (as in the Section 8 group) and an alternative way of scaling up the policy by randomly selecting recipients from a fixed eligibility pool.

Given these results, we will now explore alternative types of place-based policies that use the same financing to improve the opportunities for families who live in neighborhood C, without moving them to a different neighborhood and so, limiting the costs imposed on families living in other neighborhoods.

4.3 Place-Based Transfer Policy

We now analyze the impact effects of a “place-based transfer policy” (PBT from now on), that is, a policy that introduces a transfer to all the families living in the neighborhood with the lowest spillover.³⁴ In particular, we assume that such a policy is also unexpectedly and permanently introduced in 1990. For each scale of the MTO policy, we consider a PBT policy that levies the equivalent total amount of taxes at each time t .

Figure 4: PBT - General Equilibrium Effects



The general equilibrium effects of this policy are quite different from those generated by the MTO policy. In particular, panel (a) in Figure 4 shows the behavior of the spillover in the three neighborhoods under PBT, as a function of the policy’s scale. Under this policy, neighborhood C becomes more attractive because families living there receive a transfer. This implies that some of the poorer families living in neighborhood A will move to neighborhood C, making A more selective, and increasing A’s spillover. Panel (b) in the same figure shows that neighborhood A becomes more selective not only in terms of parental income, but also in terms of average ability of the children growing up there, given that families with more talented

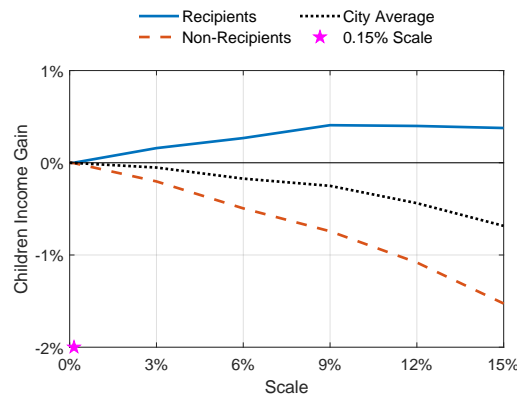
³⁴See Appendix B for more details on the introduction of the policy in the model.

children will tend not to move to C where the spillover is lower. This is very different from the MTO policy we analyzed, where the spillover in neighborhood A was deteriorating as an effect of the policy. At the same time, the spillover in neighborhood B declines because, although some of the poorer families move to C, there are fewer families moving into the neighborhood from A. This is also in contrast with what happens under the MTO policy, where neighborhood B was attracting more families from A. Finally, the spillover in neighborhood C increases, as under the MTO policy, but for different reasons: under MTO it was because the poorer families were moving out, while under PBT, it is because families from richer neighborhoods move in to receive the transfer.

Panels (c) and (d) in Figures 4 show respectively what happens to the rental rates and the sizes of the three neighborhoods as a function of the scale of the policy. The PBT policy incentivizes families to move to C. As the scale of the policy increases, families receive a larger transfer if they live in neighborhood C, and hence more families want to move there. Given the elastic housing supply, this translates in part in an increase in the rental rate of neighborhood C and in part in an increase in its size. At the same time, there is less demand for neighborhoods A and B, which translates into a reduction of both rental rates and sizes in those neighborhoods.

Next, Figure 5 shows the effects of the PBT policy on children’s future income as a function of the policy’s scale.

Figure 5: PBT - Effect on Children Income



In particular, the blue solid line shows the change in income of the children of recipient families under the PBT policy, that is, the families choosing to live in C, relative to the income of the same children in absence of the policy. The figure shows that this effect is positive, but substantially smaller than the effect for the recipient families under the MTO policy. Two forces contribute to the positive effect under PBT. First, parents receiving the transfer end up using a fraction of it, even if small, to invest in education. However, the transfer is much smaller than the housing voucher under MTO because the same finances are used to

cover a much larger group of families. Second, the spillover in C improves because of the selection of families attracted to the neighborhood by the policy. However, this effect is also quite small, especially if compared to the fact that recipient families under the MTO policy expose their children to the spillover in neighborhood A. Moreover, there is a third force dampening the positive effect. This is due to the fact that some of the recipient families who are attracted to neighborhood C to enjoy the transfer would live in neighborhoods A or B in absence of the policy, exposing their children to a higher spillover.

The red dashed line shows the effect on the income of children of non-recipient families, that is, families living in A or B, relative to the income of the same children under no policy. This effect is negative and larger the larger is the scale. This is because non-recipient families pay taxes to finance the PBT policy and so have less resources both to consume and to invest in education. This dominates the general equilibrium effect, which under PBT increases the spillover of A and B going in the opposite direction. The black dotted line shows the city-wide average of the children income change in response to the policy and shows that the overall effect is negative, although much smaller than under the MTO policy.

To sum up, the MTO policy is more effective than a PBT policy in increasing the expected income of the children in the recipient families, but it also generates larger negative effects on the other children.

4.4 Aggregate Effects

In this subsection, we explore the aggregate effects of the MTO and the PBT policy. In particular, we look at the effects of the policies on welfare, income inequality, residential segregation, and intergenerational mobility.

4.4.1 Welfare

Let us first study the welfare implications of the two policies. For simplicity, we use the utilitarian welfare criterion and assign the same weight to the utility of all parents at each point in time. Parental utility depends on the children's expected future income, but also on their own consumption and preference shocks.

The left two panels in Figure 6 refer to the welfare effects of the MTO policy and the right two panels to the welfare effects of the PBT policy.

In particular, the blue solid line in Panel (a) of Figure 6 shows the average consumption-equivalent welfare percentage gain for the recipient families under the MTO policy relative to a scenario with no policy, as a function of the policy's scale.³⁵ The first thing to notice is that the MTO policy generates a large welfare gain for the parents of recipient families. This is due both to the effect on the future income gain of their children, which we have shown in Figure 2, and to the fact that they can substantially increase their consumption, as the housing voucher now covers large part of their rents. The non-monotonicity across the scale of these welfare gains is in part due to the non-monotonicity of the children income gain that we have already

³⁵Given that the welfare gains for the recipient families are quite large, we plot their scale on the right axis.

Figure 6: Welfare Gains - MTO vs PBT



discussed, but also to the non-monotonic effect on their consumption. When the scale is very small, the recipient families are very poor and the PA policy covers their rental expenditures. However, when the scale increases the marginal recipient family is less poor and, in absence of the MTO policy, it would have to pay for rent, making the consumption effect of the MTO larger. Then, as the scale increases further, the share of income used to pay rents decreases and the effect becomes smaller again.

The dashed red line in the same panel shows the average welfare gain for the non-recipient families under the MTO program. The figure shows that on average these families suffer welfare losses. This is due both to the fact that these families are paying taxes to finance the policy and to the general equilibrium effects. In particular, families living in neighborhoods A and B have now to pay higher rents and families living in A are also exposed to a smaller spillover relative to the no-policy environment. This is evident from Panel (c) in Figure 6, which shows the average consumption-equivalent welfare gains from the MTO policy for all families depending on their birth neighborhoods and shows that families born in A and B suffer average welfare losses. Families born in C obtain average welfare gains because they include the families who receive the policy housing voucher. The share of recipient families out of the families born in C is increasing with the policy scale, making their average welfare gains monotonically increasing.

Finally, the dotted black line in panel (a) represents the average welfare gain for the city and shows that the overall effect is positive, although quite small. While average welfare gains for the recipient families are very large, the share of recipient families in the city is small and the non-recipient families suffer welfare losses dampening the overall positive effect of the policy, especially for smaller scales.

Notice that the choice of having three neighborhoods instead of two in the model has significant implications for the welfare effects of the MTO policy. In particular, the presence of a third neighborhood implies that families living in neighborhood A have the option to move to a middle neighborhood in response to recipient families moving in and reducing the spillover in A. As we show in Appendix D.4, a version of the model with only two neighborhoods tend to overestimate welfare gains from the MTO program relative to our baseline model with three neighborhoods. In response to the policy when there are only two neighborhoods, families living A either remain in A, dampening the reduction of the spillover in A, or move to C, further increasing the spillover in C. Both these effect amplify the welfare gains from the policy.

Panel (b) and (d) in Figure 6 show the analogous graphs to panel (a) and (c) for the PBT policy. The blue solid line in Panel (b) shows that the average recipient family enjoys welfare gains from the PBT policy that are increasing with the scale. However, these gains are much smaller than the ones obtained by the average recipient family under the MTO policy because under the PBT policy all families living in C receive the transfer, while under the MTO the same finances are used to give housing voucher to a much smaller fraction of the population. Moreover, the red dashed line shows that under the PBT policy, the average non-recipient family does not suffer a welfare loss like under MTO. This is because, even if they have to pay taxes to finance the policy, the general equilibrium effect for families not living in C is welfare enhancing, as rental rates both in A and B decrease and the spillover in A increases (while the one in B barely changes). This is confirmed in Panel (d), that shows that families born in all neighborhoods benefit from the policy.

To summarize the welfare implications of the two policies, let us focus on the comparison between the city-wide average consumption-equivalent welfare gain under the MTO and PBT policies, captured by the black dotted lines respectively in panels (a) and (b) of Figure 6. There are three take-aways. First, average welfare gains are larger under PBT policy, the more the larger is the scale. This is due to the fact that while welfare gains are larger for recipients under the MTO policy, the share of recipient families in the population is much smaller than the recipient families under the PBT. Moreover, under the PBT policies also non-recipient families enjoy welfare gains, while under the MTO they suffer average losses. Also, the larger the scale, the smaller are the welfare gains in terms of children's future income under MTO because of the GE effect that reduces the spillover in A and increases the cost of the policy. Second, the welfare gains for recipient families under the MTO policy are due both to the increase in children's income and to the increase in parental consumption, given that their income has to cover now only a small share of their rent. Instead, under the PBT policy most of the gains come from the increase in parental consumption, as the return of education for recipient families is not significantly affected. Third, as we emphasized in the decomposition in Figure 6, under the MTO policy the non-recipient families, who are the majority of the population, suffer welfare losses, while under the PBT policy, all families enjoy average gains. This makes

the PBT policy easier to implement, if one considers the political environment, which is outside our model.

4.4.2 Segregation, Inequality, and Intergenerational Mobility

Overall, in the previous subsection we have shown that, when we measure welfare using the utilitarian criterion, the PBT policy generates larger average welfare gains relative to the MTO policy. We now explore other aggregate implications of the two policies, such as residential segregation, income inequality, and intergenerational mobility.

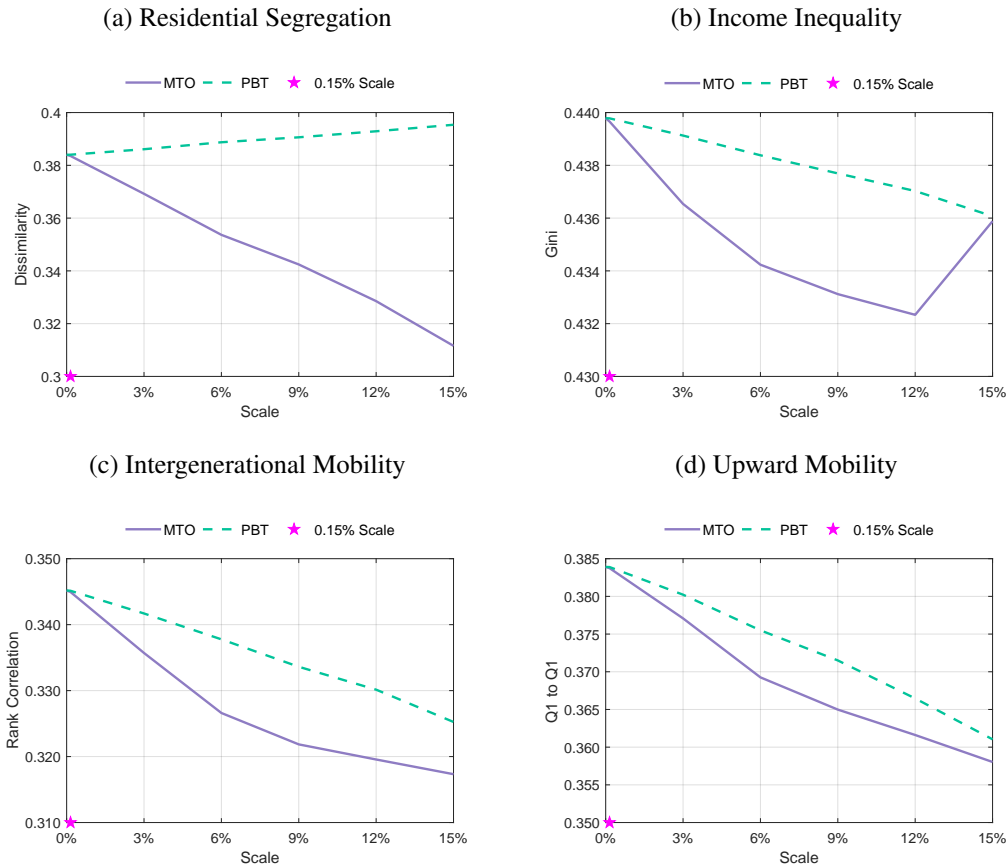
Panel (a) in Figure 7 compares the level of residential segregation by income at the time of the policy introduction as a function of the policy scale. The purple solid line refers to the MTO policy and the green dashed line to the PBT policy. When the scale of the policy is equal to the effective scale p_0 , for both policies residential segregation is at the same level as if there was no policy. However, as the scale increases, the MTO policy is successful in reducing residential segregation, while the PBT policy generates an increase in residential segregation. On the one hand, by policy design, the MTO policy reduces residential segregation by subsidizing poor families to move to neighborhood A. On the other hand, under the PBT policy the marginal families who would live in A or B with no policy have an incentive to move to C to receive the place-based transfer. This makes neighborhood A more selective and segregated. The families moving to C on average will have higher income than the neighborhood average, partly dampening the increase in segregation, but the first effect dominates.³⁶

Panel (b) in the same figure compares the level of income inequality under the two policies. The figure plots the Gini coefficient in the period in which the policy is implemented as a function of the scale of the policy. Also for inequality, when the scale of the policies is equal to p_0 , the aggregate effects of both policies are null and inequality is at the same level as if there were no policy. As the scale increases, while both policies are successful in reducing inequality, under the MTO policy the effects are larger. This is not surprising, as there is a feedback effect between residential segregation and inequality. However, the figure shows that under the MTO policy, for scales larger than $p = 12\%$, inequality increases. This is due to the fact that, as shown in Figure 1, for scales larger than $p = 12\%$, the spillover in neighborhood A declines so much that it becomes smaller than the one in neighborhood B. Given that we assume that voucher recipients have to go to the neighborhood with highest spillover in the previous period, more of the rich families sort in B and are not affected by the inflow of voucher recipients. This is why income inequality under MTO increases back up for the scale of $p = 15\%$.

Finally, panels (c) and (d) show two different measures of intergenerational mobility: the rank-rank coefficient, that represents the correlation between parental and children's income, and the Q1-to-Q1 upward mobility measure, which represents the probability that the child of a parent with income in the lowest quartile of the income distribution also will have income in the lowest quartile of the distribution. The figure shows that the MTO policy is more successful in improving both measures of intergenerational mobility

³⁶This is because neighborhood C is much larger than neighborhood A.

Figure 7: Aggregate effects



relative to the PBT policy.

Summing up, while the PBT policy generates higher welfare gains, it is not as effective as the MTO in reducing income inequality and improving intergenerational mobility and it actually increases residential segregation. We now explore a different place-based policy that could potentially resolve this tension by improving the spillover in neighborhood C and attracting families on average richer and with more talented children.

5 Place-Based Investment Policy

We now consider a place-based investment (PBI) policy that directly invests into improving the spillover of the neighborhood, i.e. investment in public school, crime reduction, information diffusion. Given that in our model, a higher spillover is beneficial for the families living in the neighborhood because it increases the returns to education, the effectiveness of this policy relies on the level of educational investment. We then assume that the government also subsidizes the cost of a basic level of education for poorer families.

In particular, the PBI policy uses the available funds for two purposes. First, the policy provides at each time

t a basic level of education \underline{e} to all the families living in neighborhood C who are in the $x - th$ percentile of the income distribution at time t . Second, the remaining resources I_t are used to directly increase the lowest spillover. As a result, the effective spillover in neighborhood n at time t , S_{nt}^* , is now

$$S_{nt}^* = S_{nt} + \psi I_{nt}^{\zeta}, \quad (6)$$

where S_{nt} is the standard spillover, as defined in equation 2, and I_{nt} is the amount of resources used to increase the spillover in neighborhood n at time t . The resources are first used to increase the spillover in neighborhood C, which is the one with the lowest spillover without policy, as long as the effective spillover in C remains lower than the one in B. If this is not the case, the resources are distributed among B and C so as to keep the effective spillovers in the two neighborhoods at the same level. If this level is higher than the effective spillover in A, then the resources are redistributed so as to keep the spillovers in the three neighborhoods equalized.³⁷

We set the basic education parameter $\underline{e} = 0.23$ to be equal to the average education in neighborhood C at the time of the policy introduction. In addition, we set the basic education eligibility at an income percentile of $x = 0.135$ to match the poverty line in the US in 1990. To further discipline the design of the PBI policy, we use the estimates about return from capital investment in public schools from Biasi et al. (2025) to pin down the parameter ζ in equation (6). They show that a cumulative \$1,000 increase in per-pupil capital expenditures raises test scores by 0.8 of a standard deviation for children in disadvantaged districts. Combining these estimates with the mapping from children test scores to their incomes as adult from Chetty et al. (2014a), we calculate that capital investment in public schools has an internal rate of return (IRR) of 4.1%. We then target the correspondent IRR in our model with a PBI policy at a scale of $p = 3\%$, which is in line with the study of Biasi et al. (2025). We obtain a calibrated value of $\zeta = 0.15$.³⁸

Panel (a) in Figure 8 shows the level of the effective spillover in the three neighborhoods at the time of the policy introduction as a function of the policy scale. By design, the policy generates convergence among the effective spillovers. In particular, at the effective scale p_0 the spillover in C increases relative to the baseline (which corresponds to the scale $p = 0$) because of the direct policy investment and the spillovers in A and B endogenously decrease because of the tax increase to finance the policy.³⁹ As the scale increases, from $p = 3\%$ onward, the funds are enough to equalize the effective spillovers in neighborhoods C and B. Panel (b) in the same figure shows that, although small, there is some convergence also in the average ability in the three neighborhoods. Panels (c) and (d) show that rents and sizes do not move much on impact.

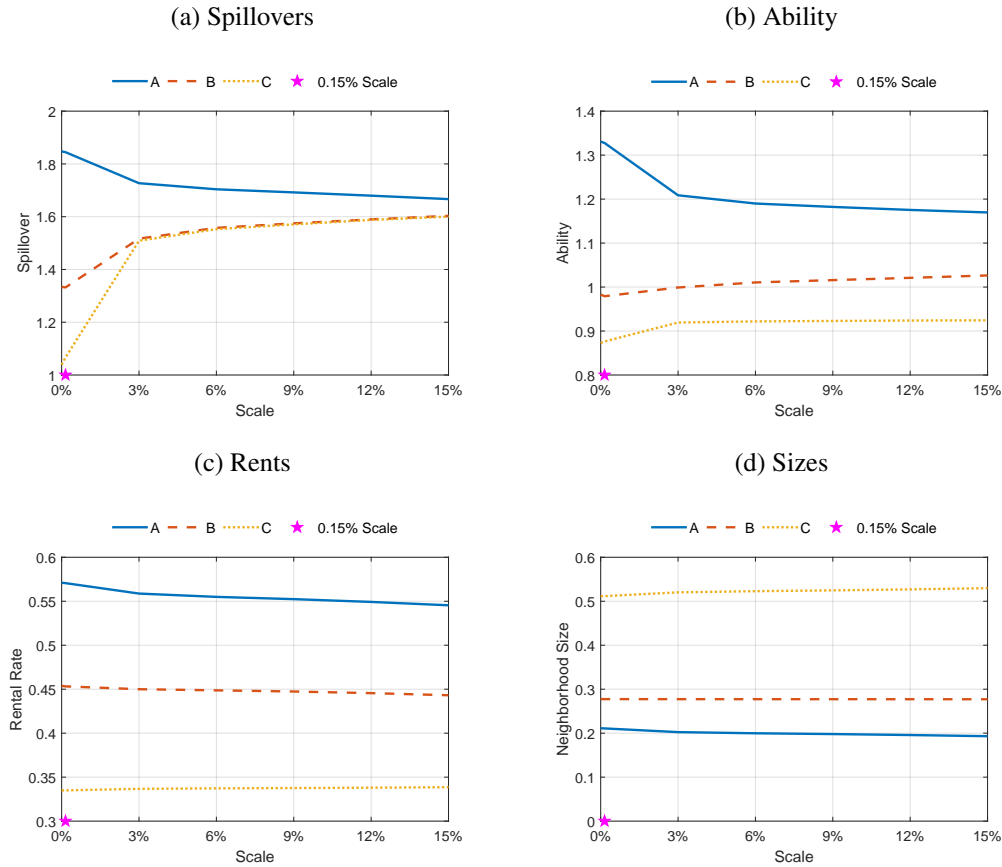
The mechanical increase of the effective spillover due to the direct policy investment is amplified by the

³⁷In particular, $I_{n_t} = \bar{I}_t$ if $S_{n_t}^* \leq S_{\tilde{n}_t}$ where \tilde{n}_t is such that $S_{n_t} < S_{\tilde{n}_t} < S_{\bar{n}_t}$. If instead, this is not true, then the total resources \bar{I}_t are split among the two neighborhoods n_t and \tilde{n}_t so that $S_{n_t}^* = S_{\tilde{n}_t}^*$, as long as $S_{n_t}^* < S_{\bar{n}_t}^*$. If this is not the case either, then the resources are split among the three neighborhoods so as to equalize the total spillovers.

³⁸See Appendix C, for details on the calculation of the IRR and on the mapping of children's test scores to income.

³⁹The neighborhood spillover is equal to the average expected children's wage as in equation 2 and we assume that the parental wage in the wage function 1 is net of taxes. This implies that as taxes increase on average in a neighborhood, there is a direct mechanical effect that decrease the spillover in that neighborhood.

Figure 8: PBI - General Equilibrium Effect



endogenous response of the families' sorting into neighborhoods. As the policy invests in increasing the spillover of a neighborhood, the opportunities for children growing up there improve and families on average richer and with higher ability children are attracted to move there. This, in turns, further increase the effective spillover because of the endogenous component. Panels (a) and (b) in Figure 9 decompose the increase in the spillover in neighborhoods B and C and show that there is an amplification effect in both neighborhoods. In particular, the solid lines represent the effective spillover with the PBI policy, while the dashed line represents the sum of the spillover with no policy (the dotted line) and the direct investment from the policy. This implies that the difference between the solid and the dashed lines represent the endogenous change in the spillover that amplifies the mechanical effect of the policy.

Let us now focus on the aggregate implications of this policy compared to the MTO and the PBT.

Panel (a) in Figure 10 compares the average welfare change in response to the PBI policy to the average welfare change under MTO and PBT that we showed in Figure 6. As the figure shows, under the PBI policy, welfare gains for small scales are negligible and become negative for larger scales, while welfare gains

Figure 9: Spillover Increase Amplification

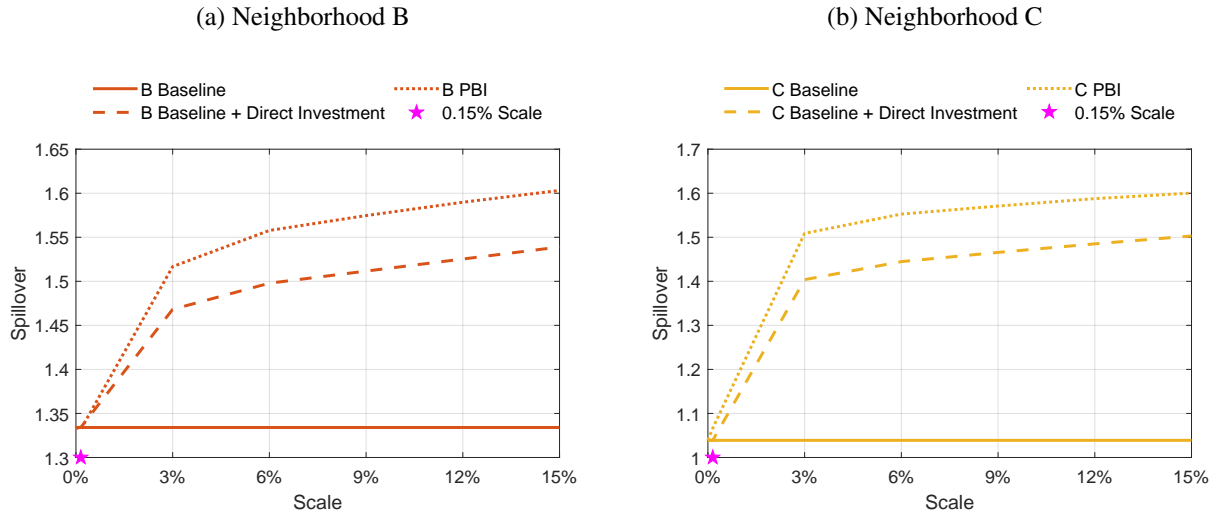
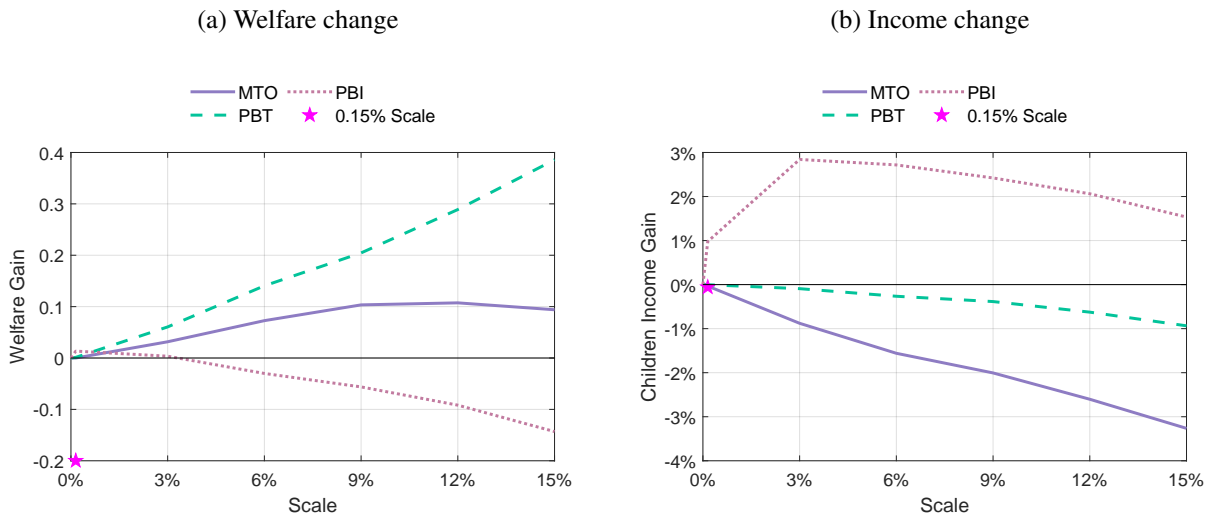


Figure 10: Welfare and Income changes across policies



are larger under MTO and, even more, under PBT. Welfare gains may come from an increase in current consumption and from an increase in children expected income. As we have already discussed, under the MTO policy, there is a large increase in expected future income for a small fraction of the population (the recipient families), but there is a decrease for the others, leading to a decline for the city average. However, there is also an increase in current consumption due to the fact that the government is now paying for a large share of the rent. Under the PBT policy, the largest component of the welfare gain comes from the consumption increase due to the transfer. The PBI policy does not transfer any resources to the recipient families, but uses the resources to increase the returns to education for the children growing up in the neighborhood with the lowest spillover, so the welfare gains could only come from gains in children's future income. However, panel (b) in the same figure shows that these gains are small if we look only at the time of the policy introduction. In the next section, we will focus on the dynamic effects of the policies, and show that the effects of the PBI become larger as time goes by and families re-sort across neighborhoods.

6 Dynamic Implications

The fact that the PBI policy on impact generates small or negative welfare effects is because, even though the policy improves the future income of children living in C, it decreases the consumption of the parents. This is because the parents do not have any transfer and instead have to pay taxes to sustain the policy. This implies that one should expect more positive effects as time goes by and the children who benefited from the better educational opportunities become parents and can increase their consumption and invest more in their children education. With this in mind, in this section, we explore the dynamic effects of the three policies and show that, with enough time, the PBI policy can resolve the tension between improving inequality and residential segregation while at the same time obtaining the largest city-wide welfare gains.

As a benchmark exercise, we fix the scale of the policy to 3%, which, as we can see in Table 3, corresponds to giving housing vouchers under the MTO to roughly 2% of the families in the population.⁴⁰

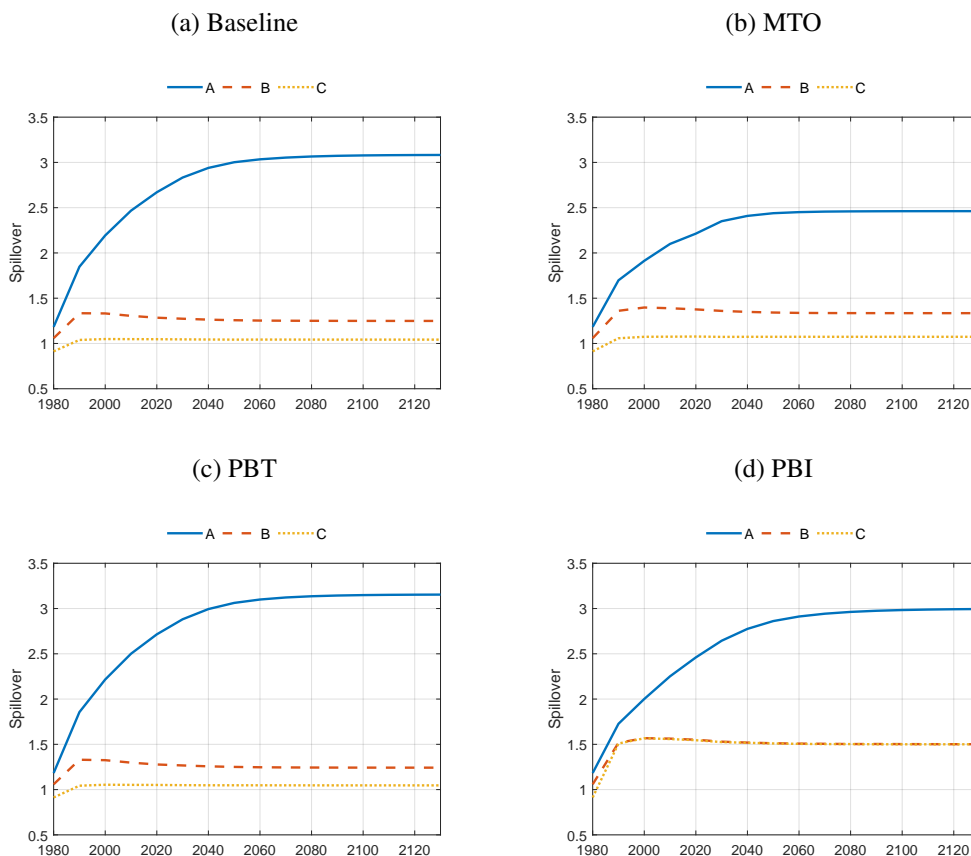
Given that the dynamics of the model capture both the effects of the policy and the effects of the underlying skill premium shock, we compare the dynamics of the model under different policy assumptions with the corresponding dynamics in the baseline model without any neighborhood-specific policy but still responding to the skill premium shock.

Figure 11 shows how the spillover in the three neighborhoods evolve over time in the baseline economy and under the three neighborhood-specific policies.

The figure shows that the PBI and MTO policies are more successful in closing the gaps between spillovers across neighborhoods. However, the MTO generates convergence mostly by decreasing the spillover in neighborhood A, while the PBI policy does so by mostly improving the spillover in C. On the contrary, the

⁴⁰We choose 3% as a benchmark because that is the scale that generates a value of investment per pupil normalized by the average wage close to the one in Biasi et al. (2025).

Figure 11: Dynamic Spillovers

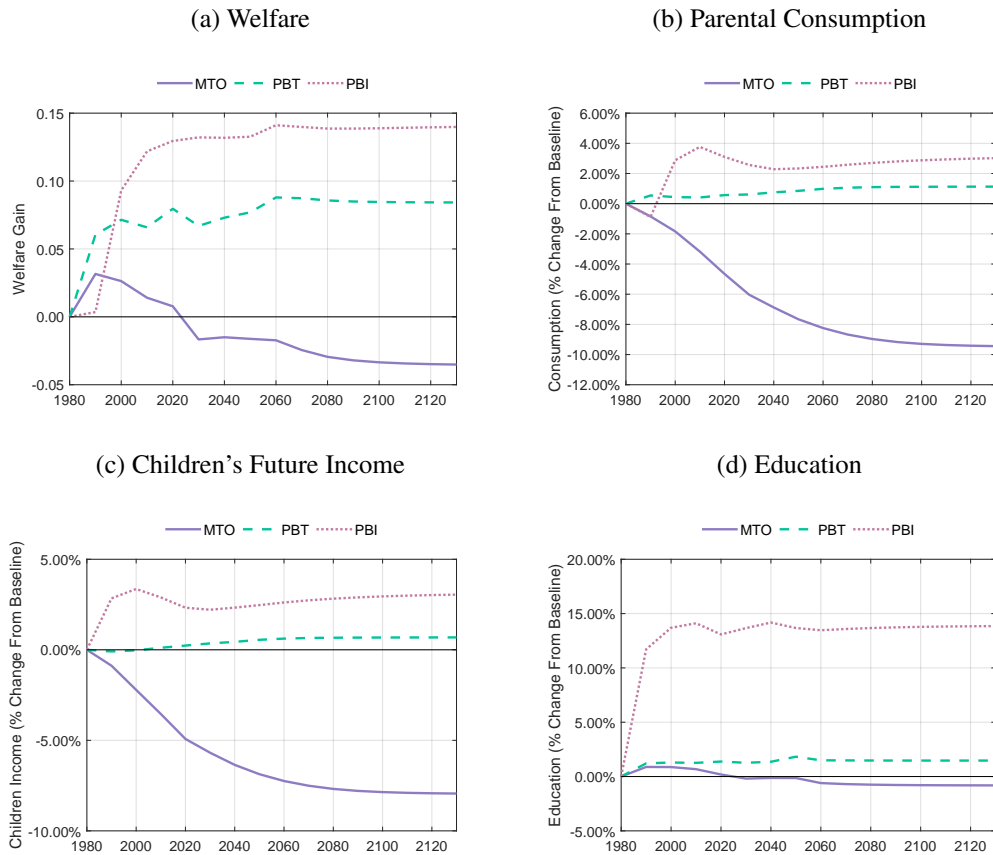


PBT policy, if anything, amplifies the gap in spillovers, especially between neighborhood A and C.

Our first main result is that, while on impact the PBI policy has smaller welfare gains relative to both the MTO and PBT policies, over time its welfare gains increase and end up dominating the other two policies.

Figure 12 explores the dynamic aggregate implications for welfare and its driving forces for the three policies. Panel (a) compares the city-wide welfare gains, in consumption equivalents, over time under the three policies. The figure confirms that on impact, that is, in 1990, the PBT policy generates the largest welfare gain and the PBI the smallest. As we have discussed, this is due to the fact that the transfer under PBT directly increases parental consumption for all the families living in neighborhood C, while the finances under PBI are purely used to invest in education (both by subsidizing a minimum level of education and by increasing the local externality), which on impact does not affect parental consumption and generates only a small increase in future children’s income. However, as time goes by, the welfare gains from the PBI policy become larger and larger. The figure shows that already one period after the policy introduction, that is, in 2000, the PBI policy generates the largest gains. This is due to the fact that the new generation

Figure 12: Dynamic Aggregate Effects



of parents have more resources for consumption and investment as they received more education and were exposed to a higher spillover when growing up under the policy. Moreover, the overall welfare under this policy keeps increasing over time as more talented kids from other neighborhoods move to *C* and parents in *C* increase the educational investment of their children to take advantage of the higher spillover, generating further increases over time in the endogenous component of the spillover. It follows that welfare increases generation after generation.⁴¹

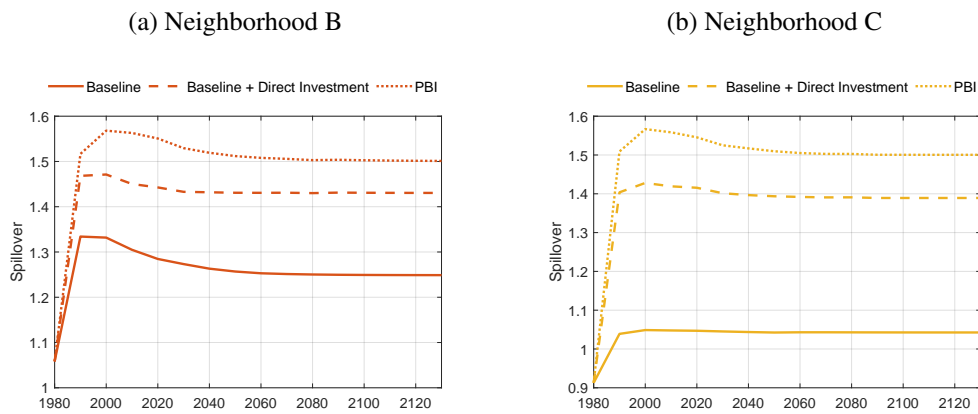
To confirm this interpretation, panels (b), (c), and (d) in the same figure respectively show city-wide averages for parental consumption, future children's income, and educational investment. Panel (b) shows that in 1990 parental consumption under the PBI policy decreases relative to the baseline because of higher taxes and because parents living in B and C use more of their resources to invest in their children's education as a response to the higher effective spillover. On the contrary, parental consumption in 1990 increases under the PBT policy thanks to the transfer to families living in C. However, starting from 2000 even parental

⁴¹In Appendix D.5, we show that the ranking of the three policies is robust to the degree of parental altruism, although the magnitudes of the welfare effects change with the relative weight on children's future income.

consumption is highest under PBI policy because of the dynamic gains received by the new generations of parents. Instead, under the MTO policy, average parental consumption is lower than the baseline and the gap increases over time. Panel (c) shows that city-wide average children’s income in 1990 is higher than in the baseline under PBI and lower under PBT and MTO. Although expected future income gains are large under the MTO for voucher receivers, the decline for the other families more than compensate for that. Finally, panel (d) shows that average education increases significantly more with PBI than under the other policies due to the higher return to education in B and C. In particular, it increases more as time goes by and new generations have benefited from higher returns in education and can, in turn, invest more in their children’s education. This channel contributes to the amplification of the welfare gains over time under PBI.

Following the same strategy of Figure 9, Figure 13 shows the increase over time in the endogenous component of the spillover in neighborhoods B and C in response to the PBI investments, for given scale of 3%. In particular, the solid lines represent the spillovers in the baseline model with no policy, the dotted lines the spillovers under PBI and the dashed lines what would be the spillover if there were only the direct investment due to the policy, without endogenous adjustment. The distance between the dotted and the dashed lines represent the endogenous response of the spillovers.

Figure 13: Spillover Dynamics

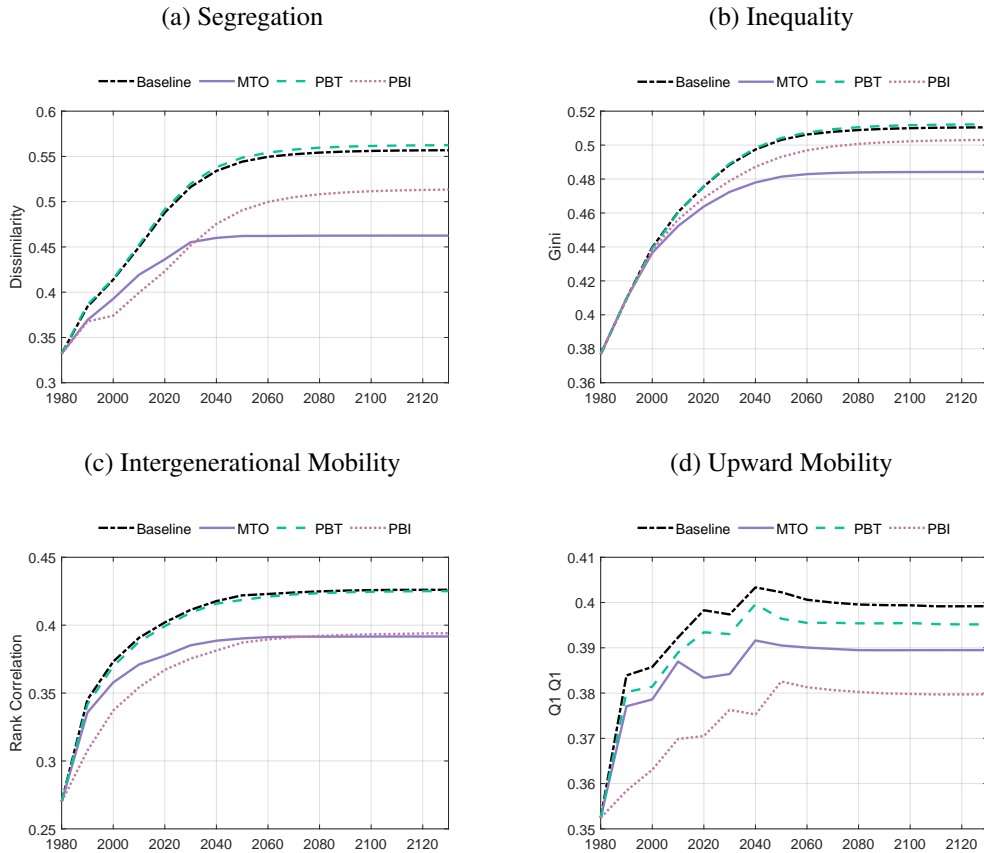


The figure shows that the endogenous component of the spillovers in both neighborhoods B and C increases over time, as families with higher income and higher ability children move to those neighborhoods in order to take advantage of the improved opportunities and all families living there invest more in education.

Figure 14 shows the dynamic implications for other aggregate outcomes. In particular, panel (a) and (b) compare the dynamic effects on residential segregation and inequality under the three different policies. The black dash-dot lines represent the baseline model with no neighborhood-specific policy and show that segregation and inequality both increase over time in response to the skill premium shock.

However, panel (a) shows that under the MTO policy, segregation increases less than in the baseline, as poor

Figure 14: Dynamic Aggregate Effects



families from neighborhood C participate in the program and use the housing voucher to move to neighborhood A. This both makes neighborhood A less selective and improves the average income in neighborhood C. However, under the PBT policy segregation increases over time more than the baseline model as the poorest families from the richer neighborhoods move to C to receive the transfer. This makes the selection of people in A even stronger and increases segregation. Finally, under the PBI policy, the spillovers in B and C are equalized, reducing the incentive to segregate by income between them and generating a level of segregation that on impact is close to the one obtained under the MTO policy, but becomes even lower over time.

Panel (b) shows that the PBT policy is the worst one also in terms of inequality dynamics, as income inequality stays over time very close to the baseline model. This is due to the fact that the transfers under the PBT policy are mostly used for parental consumption rather than to invest in children's education. On the other hand, the PBI policy, and even more, the MTO policy are more effective in reducing inequality relative to the benchmark. This is because these policy invest more in children's future income and are more effective in reducing residential segregation that contributes to amplify the reduction in inequality.

Finally, panels (c) and (d) show the two different measures of intergenerational mobility we have previously used: the rank-rank correlation and the Q1-to-Q1 transition probability. Similarly to inequality, the MTO policy is the most effective in improving intergenerational mobility also over time, while PBT is the least effective policy.

To sum up, dynamically the PBI policy resolves the tension between achieving higher welfare and improving segregation, inequality and intergenerational mobility at the same time.

6.1 Endogenous Amenities

As we have shown so far, the PBI policy dynamically increases welfare while also reducing segregation. However, in the model residential segregation is driven both by the spillovers and by the presence of exogenous amenities that are assumed to be highest in neighborhood A and lowest in neighborhood C. This means that there is limited scope to reduce residential segregation and close the gap among neighborhoods' spillovers.

It is natural to explore the implications of the policies in an economy where amenities endogenously evolve in response to the change in the distribution of families living in the different neighborhoods. In this case, the PBI policy is more effective in reducing the neighborhoods' gaps and residential segregation. In order to explore this mechanism in a simple way, we start the economy in our benchmark steady state and let amenities evolve in response to the introduction of the policy. In particular, we consider the case in which amenities depend on the average parental wage in the neighborhood, according to the following expression:

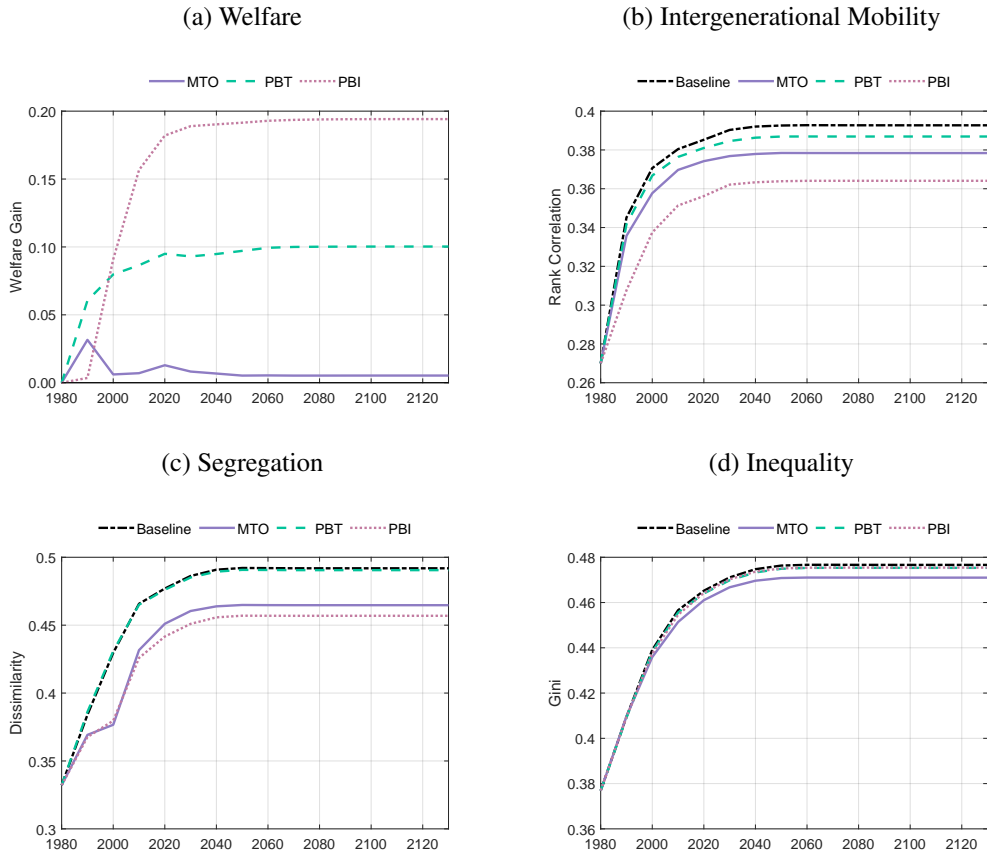
$$\theta_k = \delta_1 E[w_t | n(w_t, a_t) = k]^{\delta_2}, \quad (7)$$

for $k = A, B, C$.

Figure 15 shows the aggregate implications of such a model under the three different policies, where we choose δ_1 and δ_2 so that in steady state θ_A/θ_B and θ_C/θ_B are the same as in our benchmark model. Panel a shows that when amenities respond endogenously to the policy, the PBI policy achieves an even higher welfare gain relative to the other policies. This is because by attracting more talented children to neighborhood C, the endogenous amenities in that neighborhood increase together with the expected income of the children living there. So, families living in C not only have welfare gains because of the higher education returns, but also directly because of better amenities. This reduces the gap in both spillovers and amenities across neighborhoods. As a result, the PBI policy is now the most successful one in reducing residential segregation and improving intergenerational mobility.

In Appendix D.7, we explore the impact of a temporary version of the same three policies and show that the PBI policy is the one that has more persistent macroeconomic effects.

Figure 15: Endogenous Amenities



7 Concluding Remarks

American cities have become progressively more segregated by income over time, generating large differences in opportunities for children growing up in different neighborhoods. This trend has spurred interest in a range of neighborhood-specific policies, including the MTO program, place-based transfers (PBT), and place-based investments (PBI) in infrastructure such as public schools.

In this paper, we use a general equilibrium model with residential choice and endogenous local spillovers to examine the effects of these policies. First, we show that while housing vouchers like those in the MTO program succeed in generating income and welfare gains for recipient families, they impose negative general equilibrium effects on non-recipient families and can, at scale, even lead to welfare losses at the city level.

Using the same fiscal resources, a PBT policy—offering transfers to families living in poor neighborhoods—yields larger overall welfare gains, as families in all neighborhoods benefit either directly from the transfer or indirectly from equilibrium effects. However, this policy increases residential segregation, as poor families in richer neighborhoods relocate to poorer areas to qualify for the transfer. This dynamic, in

turn, dampens the reduction in inequality.

We finally explore a PBI policy that directly invests in local public goods to reduce the gap in spillovers across neighborhoods. This type of policy is less effective in increasing welfare in the short run, as it does not increase current resources available (as transfers and housing vouchers do). However, it generates a strong incentive to invest in the education of children, increasing resources available to the next generation and improving the spillover over time. As the spillover accumulates, welfare gains also accumulate, and the policy becomes more successful over time both in terms of increasing welfare and in terms of reducing segregation and improving the opportunities for the poor to experience the American dream.

Appendix

A Details on the MTO Policy

The MTO program is a randomized experiment conducted by the US Department of Housing and Urban Development (HUD) to improve opportunities for children of low-income families living in poor neighborhoods. This program run from 1994 to 1998 in five cities: Baltimore, Boston, Chicago, Los Angeles, and New York.

To be eligible for the program, families had to satisfy a number of criteria: 1) they needed to have children below 18 years old of age; 2) they had to live in high-poverty neighborhoods (with more than 40% of population below the poverty line); 3) they had to live in public housing developments or project-based assisted housing, which required them to be low income families (below 50% of the local median).

The public housing authority (PHA) in each city reached out to all eligible families and gave them the opportunity to pre-apply for the program. After a preliminary eligibility screening, families were randomly ordered into a waiting list. Then, small groups of families from the top of the list were called to attend a visit at the PHA, where they were told the rules of the experiment and then they had to decide whether they wished to go forward with the application. If they decided to apply, they were randomly assigned to one of three groups:

1. the experimental group was given housing vouchers that could be used only to move to census tracts with 1990 poverty rates below 10%;
2. the Section 8 group received regular Section 8 housing vouchers that could be used without any specific relocation constraint;
3. the control group received no assistance through the MTO program, but continued to be eligible for housing assistance and other welfare programs.

Voucher recipients were required to contribute 30% of their annual household income towards rent and utilities and received housing vouchers that covered the difference between their rent and the family's contribution, up to a maximum amount, defined as the 40th percentile of rental costs in a metro area (Fair Market Rent). Families remained eligible for these vouchers indefinitely as long as their income was below 50 percent of the median income in their metro area.

Table A1 reports site-level eligible, enrolled, and compliance counts. The first two columns report for each site, respectively, the total population in 1990 and the estimated number of eligible families, as reported in Feins et al. (1996). The table also reports the number of families that were effectively enrolled in the MTO program according to Sanbonmatsu et al. (2011) and their assignment to the three different groups. The enrolled families are a bit less than one third of the eligible families. This is because, given the total finances

available and the estimate of compliance rates, the program was designed to target an enrollment of 4,436 families, which is roughly the number of families who ended up being enrolled. As we explained above the criterion to select the families from the waiting list was random. However, there was selection in the families who effectively pre-applied for the program and also for the ones who decided to apply after their visit to the PHA. Moreover, one of the eligibility criterion was to be eligible for the Section 8 program. While the main criterion for Section 8 eligibility is to be a "low-income" family (that is, with income not higher than 80% of the median), the program explicitly targets poorer families. Roughly one third of eligible families are ultimately enrolled (4,142 of 15,545), consistent with program targets (Feins et al., 1996; Sanbonmatsu et al., 2011).⁴²

Finally, the table also shows the compliance rate, which is the share of enrolled families in the experimental and section 8 group that actually used the voucher. As further investigated by Bergman et al. (forthcoming), many frictions might have prevented some families to actually use the voucher, including housing availability, landlords' skimming, liquidity constraints. The compliance rate is higher for the section 8 group, as the families were not constrained to find housing in a restricted set of neighborhoods.

Table A1: Eligible, enrolled, and compliant families

MTO Site	Population in 1990	All Groups		Experimental Group		Section 8 Group	
		Eligible Families	Enrolled Families	Enrolled Families	Compliance Rate	Enrolled Families	Compliance Rate
Baltimore	736,014	2,300	572	252	53.50%	123	79.80%
Boston	574,289	4,500	868	366	43.60%	176	51.10%
Chicago	2,783,660	2,415	825	460	33.40%	133	67.40%
Los Angeles	3,485,499	3,900	929	340	60.50%	200	71.60%
New York	7,322,564	2,430	948	401	46.40%	252	45.20%
All Sites	13,902,026	15,545	4,142	1,819	47.40%	844	61.60%

Initial research evaluating the effects of the MTO program (e.g. Katz et al. (2001), Kling et al. (2007), Clampet-Lundquist and Massey (2008), Ludwig et al. (2013)) did not find significant impact of housing vouchers on future income and employment outcomes of children of recipient families. They found only effects on non-economic measures of well-being, such as physical and mental health.

However, Chetty et al. (2016) have recently revisited the economic impact of the MTO program, using more recent administrative data that include adult labor market outcomes of individuals who were young children when their families participated in the program. Motivated by their previous work on causal effects of neighborhood exposure, they split the sample of families who participated in the MTO program into

⁴²From the website <https://www.cbpp.org/research/policy-basics-section-8-project-based-rental-assistance>: "For each participating housing development, at least 40 percent of the subsidized units that become available annually must go to families with "extremely low incomes" (up to the poverty line or 30 percent of the local median, whichever is higher). Most of the remaining units are restricted to families or individuals with incomes not above half of the local median."

two groups: those with children younger than 13 years old at the time of the program and those with older children. With this distinction, they find significant economic effects for children who moved to lower poverty neighborhoods. In particular, they look at the treatment-on-the-treated (TOT) effect, which is the change in adult income for children of voucher recipient families who used the voucher relative to the control group. For children who were younger than 13 when their family took up an experimental voucher, they estimate an increase in adult annual income by the time they reached their mid-twenties of \$3,477, which corresponds to a 31% increase relative to a mean of \$11,270 in the control group. For children younger than 13 of families in the Section 8 group, they estimate an increase in income of 15% relative to the control group. Instead, they find negligible or slightly negative effects for the children who were older than 13 years old when their families used the voucher. They also find significant effects on college attendance for children younger than 13.

B Mapping the Policies to the Model

We now describe how we map the three types of neighborhood-specific policies to the model.

B.1 MTO Policy

The first policy that we consider is a housing voucher policy that mimic the MTO program implemented in the US. We assume that the policy is permanent and is unexpectedly introduced at time \bar{t} . For each $t \geq \bar{t}$, let $\chi(w_t, a_t, b_t)$ denote the eligibility indicator, with $\chi_t(w_t, a_t, b_t) = 1$ if the parent is eligible for the program and equal to 0 otherwise. To map the MTO policy eligibility criteria, we require that families need to: 1) live in the neighborhood with the lowest spillover, and 2) belong to the poorest p -th percentile of the metro income distribution.⁴³ That is, we assume that $\chi(w_t, a_t, b_t) = 1$ if $b_t = \underline{n}_t$ and $w_t \leq \tilde{w}_t$, where \tilde{w}_t is such that

$$F_w[\tilde{w}_t] = p.$$

In our model the neighborhood with the lowest spillover is also the one where the poorer families live. However, given that local spillovers are endogenous and evolve over time, and so does the composition of the neighborhood, the worst neighborhood may change in response to the policy.

Eligible families are offered a housing voucher to cover the difference between their rent and the family's

⁴³In our model, we do not distinguish between eligible and enrolled families and we simply assume a tighter wage cutoff for enrollment (instead of selecting a random subset of the enrolled families), given that the documentation suggests that the poorest of the eligible families were the ones eventually enrolled in the program. In particular, there are three reasons why this seems to be the case: 1) the tracts effectively targeted by the program had average poverty rate of 57% (see Appendix B in the MTO Congress Report), which is stricter than what the eligibility criterion would have required; 2) one of the criterion to be enrolled for the MTO program was to live in public housing developments or project-based assisted housing, which give priority to the poorer families; 3) the families who enrolled in the MTO program were poorer than the families living in the same public housing developments who did not enroll in the MTO (see Table 5 from Goering et al. (1999)).

contribution, equal to a fraction q of their income, with a cap \bar{r} . We focus on the “experimental group” and assume that families accepting the voucher have to move to the neighborhood with the highest spillover, \bar{n}_t . We assume that the policy is financed with income taxes on the whole population and that the government budget has to balance. The public assistance program is still in place and if a family is eligible for both, she will get the housing voucher transfer from the MTO program, with the same contribution rules, and the same consumption transfer $T_t(w_t)$ from the public assistance program. The eligibility cut-off for the public assistance program is unchanged, so that families who were eligible without the MTO program in place, are still eligible afterwards.⁴⁴ The government budget balance condition can be written as

$$\int_{v_t=1} \min\{(R_{\bar{n}_t} - qw_t), \bar{r}\} dF_t(w_t, a_t, b_t) + \int_{w_t < \hat{w}_t \& v_t=0} (R_{\bar{n}_t} + \underline{c} - (1 - \kappa_t)w_t) dF_t(w_t, a_t, b_t) \\ \int_{w_t < \hat{w}_t \& v_t=1} [\underline{c} - \max\{(1 - \kappa_t)w_t - R_{\bar{n}_t}, 0\}] dF_t(w_t, a_t, b_t) \leq \kappa_t \left(\int w_t dF_t(w_t, a_t, b_t) \right).$$

Parents who are eligible for the program take up the voucher if the utility from using it is larger than the one from not using it. Define $v_t(w_t, a_t, b_t)$ as the voucher take-up indicator. Let us first consider the case of experimental policy. For all parents such that $\chi_t(w_t, a_t, b_t) = 0$, $v_t(w_t, a_t, b_t) = 0$, while for all parents such that $\chi_t(w_t, a_t, b_t) = 1$ with $w_t \geq \hat{w}_t$, $v_t(w_t, a_t, b_t)$ solves

$$\max_{v_t} \{U^V(w_t, a_t, b_t), U^N(w_t, a_t, b_t)\}, \quad (8)$$

where

$$U^V(w_t, a_t, b_t) = \max_{e_t} \log[\theta_{\bar{n}_t}((1 - \kappa_t)w_t - \max\{qw_t, R_{\bar{n}_t} - \bar{r}\} + s_t(w_t)\Pi_t I_{w_t \geq \bar{w}_t} - \tau e_t^\gamma)] \\ + \log[\Omega(w_t, a_t, e_t, S_{\bar{n}_t}, \varepsilon_t)] + \sigma_\zeta \zeta_{\bar{n}_t} - \mu I_{n_t \neq b_t}, \quad (9)$$

and

$$U^N(w_t, a_t, b_t) = \max_{e_t, n_t} \log[\theta_{n_t}((1 - \kappa_t)w_t + s_t(w_t)\Pi_t I_{w_t \geq \bar{w}_t} - R_{n_t} - \tau e_t^\gamma)] + \\ + \log[\Omega(w_t, a_t, e_t, S_{n_t}, \varepsilon_t)] + \sigma_\zeta \zeta_{n_t} - \mu I_{n_t \neq b_t}, \quad (10)$$

where $U^V(w_t, a_t, b_t)$ is the value of taking up the voucher, $U^N(w_t, a_t, b_t)$ the value of not taking it up, and $\Omega(w_t, a_t, e_t, S_{n_t}, \varepsilon_t) = (y + a_t e_t \eta (\beta_0 + \beta_1 S_{n_t})^\xi) w_t^\alpha \varepsilon_t$. A parent taking up the voucher has to move to the neighborhood with the highest spillover and has to use a fraction q of her income to contribute for the rent. The voucher covers the difference between $R_{\hat{n}_t}$ and her contribution up to the cap \bar{r} . A parent not taking up the voucher has to pay the full rent, pay taxes on the total income, but can choose the optimal neighborhood.

⁴⁴Notice that families who are eligible for both the MTO and the PA programs, may end up consuming more than \underline{c} .

For a parent with $w_t < \hat{w}_t$, the value of taking up the voucher is the same as in (9), except that she will get a transfer $T_t(w_t) = \underline{c} - \max\{(1 - \kappa_t)w_t - R_{n_t}, 0\}$ that she is forced to consume. Moreover, the value of not taking it up is equal to (10), except that she now is subject to the standard rules of the PA program, that is, she has to choose neighborhood n_t , she has to use her income to pay as much as she can of the rent, and she gets a transfer $T_t(w_t) = \underline{c} - \max\{(1 - \kappa_t)w_t - R_{n_t}, 0\}$ that needs to be consumed.⁴⁵

There are two main reasons why the model can generate a take-up rate smaller than 100%. First, the rental rate in the worse neighborhood may be lower than the required down-payment to move to the best neighborhood, as required if accepting the voucher. Moreover, the required down-payment may be higher than 30% of the income if the voucher cap is binding. Second, the moving costs together with the idiosyncratic preference shocks may be such that some eligible families may prefer to remain in their original neighborhood.

B.2 PBT Policy

We now introduce the “place-based transfer policy” (PBT from now on), that is, a policy that introduces a transfer to all the families living in the poorest neighborhood n_t . In particular, we assume that such a policy is also unexpectedly introduced at time \bar{t} , and gives at any time $t \geq \bar{t}$ a transfer \tilde{T}_t to all parents choosing to live in the neighborhood with the lowest spillover n_t . As for the case of the MTO, the public assistance program is still in place and if a family is eligible for both, she will get the transfer \tilde{T}_t on top of the PA transfer $T_t(w_t) = \underline{c} - \max\{(1 - \kappa_t)w_t - R_{n_t}, 0\}$. As before, the government finance both the PA and the PBT programs with income taxes and the budget balance condition at time t is now given by

$$\begin{aligned} & \int_{w_t < \hat{w}_t} (R_{n_t} + \underline{c} - (1 - \kappa_t)w_t) dF_t(w_t, a_t, b_t) + \tilde{T}_t \left(\int_{n_t = n_t} dF_t(w_t, a_t, b_t) \right) \\ & \leq \kappa_t \left(\int w_t dF_t(w_t, a_t, b_t) \right). \end{aligned}$$

To compare the effects of the two different policies, for each scale p of the MTO program, we pick the transfer such that the total taxes collected are the same as the ones collected with the MTO. First, we find the tax rate κ_t such that the total tax revenues used under PBT are equal to the total tax revenues used under the MTO program for each scale p .⁴⁶ We then back up as a residual the place-based transfer \tilde{T}_t . However, given that parents make their residential choice knowing that the policy is in place, this becomes a fixed point problem.

B.3 PBI Policy

We now consider an alternative type of place-based policy, where the same finances are used to directly invest into improving the spillover of the neighborhood, i.e. investment in public school, crime reduction,

⁴⁵The equilibrium definition is the natural generalization of the equilibrium defined in section 2.2.

⁴⁶At time \bar{t} , this means simply that we keep the tax rate κ_t the same as under MTO. However, for the following periods $t > \bar{t}$, the tax rate might be different because the distribution $F_t(w_t, a_t, b_t)$ evolves endogenously in a different way under the two policies.

information diffusion. Given that in this model, a higher spillover is beneficial for the families living in the neighborhood because it increases the returns to education, the effectiveness of this policy relies on parents investing in education. We then assume that the government also subsidizes the cost of some basic level of education for those families that are so poor that would not have otherwise afforded it.

As we did for the analysis of the PBT, for each scale of the MTO policy p , we compare a place-based investment policy that uses the same total finances. The funds are used to pay for a basic level of education \underline{e} for families living in n_t that belong to the lowest x -th percentile of the city-wide income distribution. Let us define \hat{w}_t the cut-off such that $F(\hat{w}_t) = x\%$. The remaining resources \bar{I}_t are used to increase directly the spillover of neighborhood n_t , and, once it reaches the level of the second highest spillover, to further increase both the second highest and the lowest spillovers. The government budget balance condition is now

$$\int_{w_t < \hat{w}_t} (R_{n,t} + \underline{c} - (1 - \kappa_t)w_t) dF_t(w_t, a_t, b_t) + \tau e^\gamma \left(\int_{n_t = n_t, w_t < \hat{w}_t} dF_t(w_t, a_t, b_t) \right) + \bar{I}_t \leq \kappa_t \left(\int w_t dF_t(w_t, a_t, b_t) \right),$$

where the right-hand-side is the total level of finances equalized to the finances used under MTO for each scale p , and \bar{I}_t is the residual amount of resources after paying for basic education and PA programs. The effective spillover in neighborhood n at time t , S_{nt}^* , is now

$$S_{nt}^* = S_{nt} + \psi I_{nt}^\zeta, \quad (11)$$

where S_{nt} is the standard spillover, as defined in equation 2, and I_{nt} is the amount of resources used to increase the spillover in neighborhood n at time t . In particular, $I_{n,t} = \bar{I}_t$ if $S_{n,t}^* \leq S_{\tilde{n},t}$ where \tilde{n}_t is such that $S_{n,t} < S_{\tilde{n},t} < S_{\bar{n},t}$. If instead, this is not true, then the total resources \bar{I}_t are split among the two neighborhoods n_t and \tilde{n}_t so that $S_{n,t}^* = S_{\tilde{n},t}^*$, as long as $S_{n,t}^* < S_{\bar{n},t}^*$. If this is not the case either, then the resources are split among the three neighborhoods so as to equalize the total spillovers.

C Additional Details on the Calibration

C.1 Normalizations

For convenience, let us report the optimization problem for a household with wage w and a child with latent productivity a ,

$$U(w_t, a_t, b_t) = \max_{c_t, e_t, n_t} \log(\theta_{n_t} c_t) + E_t[\log(w_{t+1})] + \sigma_\zeta \zeta_{n_t} - \mu I_{n_t \neq b_t}$$

and her consumption and future wage,

$$c_t = (1 - \kappa_t)w_t + s_t(w_t) \Pi_t I_{w_t > \bar{w}_t} - R_{n_t} - \tau e_t^\gamma$$

$$w_{t+1} = (y + a_t e_t \eta_t (\beta_0 + \beta_1 S_{n_t})^\xi) w_t^\alpha \varepsilon_t$$

First note that average latent productivity is not independent of β_0 and β_1 , as we can scale a by a constant c_a and scale both β_0 and β_1 by $c_a^{-\frac{1}{\xi}}$ while leaving the optimization problem and the wage expression unchanged. Specifically, we can set $c_a = \frac{1}{\mu_a}$ so that the adjusted average latent productivity is equal to 1.

Moreover, we can scale ε by a constant c_ε and, at the same time, scale y by c_ε^{-1} and both β_0 and β_1 by $c_\varepsilon^{-\frac{1}{\xi}}$, again leaving the problem unchanged. We can normalize the mean of ε to 1 by setting $c_\varepsilon = \frac{1}{\mu_\varepsilon}$.

We also note that the value of λ_C is pinned down by the normalization of the city size to 1 in 1980.

Next, notice that we can multiply w by a constant c_w . We can scale y by $c_w^{-(1-\alpha)}$ and β_0 and β_1 by $c_w^{-\frac{(1-\alpha)}{\xi}}$. This leaves the wage dynamics unchanged. Moreover, from the housing market condition, R_n is going to be automatically scaled up by the same constant, and we can multiply τ by c_w so that the optimization problem is unchanged as well. This means that we can choose $c_w > 0$ so that the average wage in the economy is equal to 1 in 1980.

Finally, we show that we can normalize τ . In particular, we can make the transformation $\tilde{e} = \tau^{\frac{1}{\gamma}} e$ and scale β_0 and β_1 by $\tau^{\frac{1}{\gamma\xi}}$. This leaves the optimization problem (where we now optimize over n and \tilde{e} instead of n and e) and the wage equation unchanged.

C.2 Population Growth

To calibrate to population growth rate, we rely on the statistics and projections from the CBO’s January 2025 report “The Demographic Outlook: 2025 to 2055”⁴⁷. In the second column of Table C2 we report the population growth by decade from 1980 to 2060 using the CBO data up to 2024 and the CBO projections from 2025 onward.

Since our calibration focuses on a sample of metropolitan areas, we adjust the CBO population growth figures, which represent the entire U.S. population. Between 1980 and 2010, the average population growth in our metro sample, based on Census data, is 8.9%, compared to 10.3% for the overall U.S. population according to the CBO. To align these measures, we rescale the CBO growth rates by a factor of 0.9, which is the ratio of metro population growth to total population growth, so that the adjusted values match the metro growth rate over 1980–2010. This same rescaling factor is applied to all projected values through 2060, resulting in the growth rates shown in the third column of Table C2.

Because population growth is projected to reach 0% by 2050, we assume a 0% growth rate for all subsequent decades until the model converges to a new steady state.

C.3 Calibration of the PBI parameters

The calibration of the parameters for the PBI policy is based on the following steps. First, we use the estimates on the return from capital investment in public schools from Biasi et al. (2025). Combining these

⁴⁷<https://www.cbo.gov/system/files/2025-01/60875-demographic-outlook.pdf>

Table C2: CBO Population Growth and Rescaled Growth by MSA Population (1980–2060)

Period	CBO Population Growth	Rescaled by MSA Growth Ratio
1980–1990	10.2%	8.8%
1990–2000	11.0%	9.5%
2000–2010	9.7%	8.4%
2010–2020	7.3%	6.3%
2020–2030	6.5%	5.6%
2030–2040	2.6%	2.3%
2040–2050	1.0%	0.8%
2050–2060	0.0%	0.0%

estimates with the mapping from children test scores to their incomes as adult from Chetty et al. (2014a), we calculate the internal rate of return (IRR) of capital investment in public schools. We then target the correspondent IRR in our model with a PBI policy at a scale of $p = 3\%$, which is in line with the study of Biasi et al. (2025). We obtain a calibrated value of $\zeta = 0.15$. Below we illustrate this procedure in detail.

Biasi et al. (2025) show that a cumulative \$1,000 increase in per-pupil capital expenditures raises test scores by 0.8 of a standard deviation (SD) for children in disadvantaged districts, using data for the period 1995–2017. Disadvantaged districts are identified by grouping them according to the share of students eligible for free or reduced-price meals, a proxy for low socio-economic status (SES). Their analysis focuses on districts in the bottom and top terciles of the distribution of this share across all U.S. districts in 1995.

Next, we map the increase in children test scores into income gains as adult using the estimates from Chetty et al. (2014a). They use data on test scores for English language, arts, and math for students in grades 3–8 in every year from the spring of 1989 to 2009 (with the exception of 7th grade English scores in 2002) and they measure adult outcomes for the years 1996–2011. Their findings indicate that a 1 SD increase in the current test score is associated with \$2,600 increase in earnings on average at age 28, which corresponds to a 12% income gain relative to the average earnings in the sample of \$ 20,885. We note that the average earnings in the sample are similar to the average wage in the US in 1990, which was about \$21,000⁴⁸.

We then calculate the annual earnings gain from the public school capital investment by assuming that the intervention raises test scores by 0.08 SD, and that each 1 SD increase in test scores leads to a 12% increase in income. Applying this, the annual gain is approximately $0.08 \times 0.12 \times \$21,000 = 201.60$.

We can then calculate the internal rate of return (IRR) of the public schools capital investment in two ways, either (i) calculating the Net Present Value (NPV) of the income gain over an individual’s working life, or (ii) using a simplified two-period IRR similar to the structure of our model.

Working Life NPV Calculation. We calculate the NPV of the public school capital investment, assuming that the cost of \$1,000 is incurred over ten years while a child is in school, and that a gain of \$201.60 is

⁴⁸Data on the US National Average Wage Index from the SSA: <https://www.ssa.gov/oact/cola/AWI.html>

obtained for each year of the individual’s adult working life, from age 25 to 55. This corresponds to the following investment streams: the cost of the investment is \$100 per year when the child is aged 8 to 17. Next, the income stream is zero when the individual is aged 18–24, before entering the adult workforce. Finally, the gain from the investment is \$201.60 per year when the individual is aged 25 to 55. The IRR of the investment satisfies the following equation:

$$\sum_{t=8}^{55} \left(\frac{1}{1 + IRR} \right)^t x_t = 0$$

where x_t indicates the investment flows, which are $x_t = -100$ for $t \in [8, 17]$, $x_t = 0$ for $t \in [18, 24]$, and $x_t = 201.60$ for $t \in [25, 55]$. This formula delivers an IRR of 7.6%.

Two-period IRR calculation. Alternatively, we consider an IRR calculation over two time periods, following the structure of our model, where individuals are children for one period, and then adults for one period, with each period lasting 10 years. We then assume that the cost of \$1,000 is incurred in the first period, when the individual is a child. The gain accrues in the second period, when the individual is an adult, and earnings increase by \$201.60 for each of the ten years of the second period. We then calculate the gain of the program over 10 years as $\$201.60 \times 10 = \$2,016$ and divide it by the cost over 10 years, using the following formula:

$$IRR = \left(\frac{\text{return over 10 years}}{\text{cost over 10 years}} \right)^{1/10} - 1 = \left(\frac{10 \times 0.08 \times 0.12 \times \$21,000}{\$1,000} \right)^{1/10} - 1 = 7.3\% \quad (12)$$

We note that the two methodologies deliver a remarkably similar IRR for the program.

As a last step, we implement the PBI in our model and use the IRR as a target to calibrate the parameter ζ . We pick the more conservative IRR target of 7.3% from the second methodology.

To implement the PBI in our model, we need to first determine the transfer per pupil. To map the \$1,000 capital investment from Biasi et al. (2025) into our model, we benchmark it to the average wage. We consider the national average wage between 1990 and 2020, which approximately span their sample period, which was approximately \$ 38,000. The expenditure per pupil could then be expressed as a fraction of the average wage, equal to $\$1,000 / \$ 38,000 = 0.026$. We then implement the PBI in our model in 1990 with a transfer per pupil that amounts to 2.6% of the average 1990 wage within our model. We also note that this transfer generates a tax rate that approximately corresponds to the tax rate of the MTO at 3% scale.

Finally, we calculate the IRR of the PBI in our model using the formula in equation (12). To do this, we calculate the income gain for the recipient children in the year 2000 (one period after the policy implementation.) Given that the results of Biasi et al. (2025) apply to disadvantaged children in the bottom tercile of districts, we calculate the income for children in the bottom tercile of the income distribution. Targeting an IRR of 7.3%, we obtain a value of $\zeta = 0.15$.

D Additional Results

D.1 Alternative MTO Specifications

In this section, we explore alternative ways of implementing the MTO policy, by changing the destination neighborhood required to obtain the housing voucher. We do that with two purposes. First, it is not obvious which neighborhood in the model better represents the “low poverty neighborhood” where voucher recipients need to move according to the MTO program if they are in the experimental group. In our baseline exercise, we assumed that this was neighborhood A, while here we consider the alternative exercise where voucher recipients have to move to neighborhood B. Second, some of the families participating in the MTO program were assigned to the “section 8” group, which allows them to choose the neighborhood where to move. We then consider also the exercise where voucher receivers can choose whether to move to neighborhood A or B to represent that case.

Figure D1 presents the results for these two alternative specifications of the MTO policy. The solid line represents our baseline specification. The dashed line represents the case where voucher recipients are required to move to neighborhood B.⁴⁹ The dotted line represents a case where voucher recipients can choose whether to move to neighborhood A or B.

Panel (a) displays the effects on the expected income gain for the children of recipient families. When recipients move to neighborhood B rather than A, the expected income gain for the children is lower, given that the spillover in B is lower than in A. When recipients are allowed to choose whether to go to A or B, most of them choose B at small scales, while roughly half of them choose B at larger scales.⁵⁰ This is because as the scale increases, there are more rich families among voucher recipients, who can afford the neighborhood with highest spillover. As a result, the income gain is closer to the case where all recipients go to B for small scales, and it becomes more similar to the case where all recipients go to A as the scale increases. Note that, at the largest scales, the income gain is even higher than in our baseline specification. The reason is that, when recipients go partly to A and partly to B, the general equilibrium effect that reduces the spillover in A is lower than in the baseline. As a consequence, the gain for the children of the recipients that choose A is larger.

Panels (b), (c), and (d) report the aggregate effects on income inequality, residential segregation, and intergenerational mobility. We observe that the case where recipients move to A is the most effective in improving these aggregate outcomes, given that individuals are required to move to the neighborhood with the highest spillover. The effect is slightly lower for the case of moving to either A or B, and even lower for the case of moving to B, where recipients are exposed to a lower spillover.

⁴⁹For this case, we set the voucher cup $\bar{r} = .8 * R_{B,1990}$.

⁵⁰The fraction of recipients choosing neighborhood B, conditional on taking up the voucher, is 99.9% at 0.15% scale, 63.5% at 3% scale, 55.9% at 6% scale, 52.9% at 9% scale, 50.1% at 12% scale, and 47.7% at 15% scale.

Figure D1: Alternative MTO Specifications



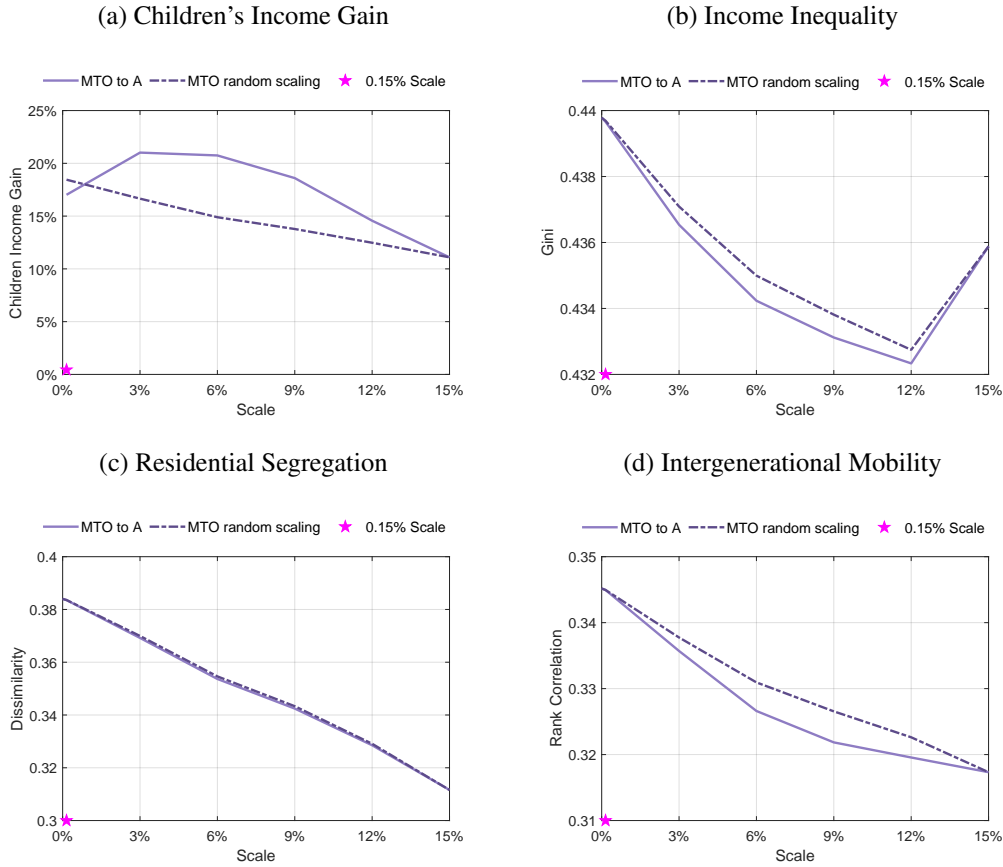
D.2 Alternative Scaling Up Specification

In this section, we consider an alternative way of scaling up the MTO policy, by fixing the eligibility pool and increasing the share of families that are offered the housing voucher. Specifically, we fix the income percentile for eligibility at 15% across all scales. The voucher is then offered to a random subset of families, with mass x , within the eligible group. The scaling up is implemented by increasing the mass x of families receiving the voucher, while keeping the eligible group fixed. At each scale, we choose the mass x so as to ensure the same number of families are offered the voucher as in our baseline specification.

Figure D2 compares the results from our baseline specification (solid) line, and the alternative method of scaling up the policy (dash-dotted line).

Panel (a) reports the expected income gain for the children of recipient families. The income gain for the case of scaling up targeting a random share of a fixed eligible pool is monotonically decreasing, unlike the baseline case. The reason is that, in the baseline, scaling up consists in both increasing the share of eligible individuals and changing their composition. The composition of eligible individual changes because the income percentile for eligibility increases, so eligible individuals have higher income at larger scales; in addition, the average ability of eligible children also increases, given the correlation between income and

Figure D2: Alternative MTO Specifications



ability and the inter-generational correlation of ability in our model. The changing composition of the eligible pool drives the non-monotonicity in the income gain. On the contrary, in the case where the eligible pool is fixed and recipients are randomly selected, the income and ability composition of recipients are fixed. As a result, the decrease in income gain as the scale increases is driven by general equilibrium effects of the policy, that reduce the spillover in A and increase the rent in A, reducing the benefit of moving for recipients.

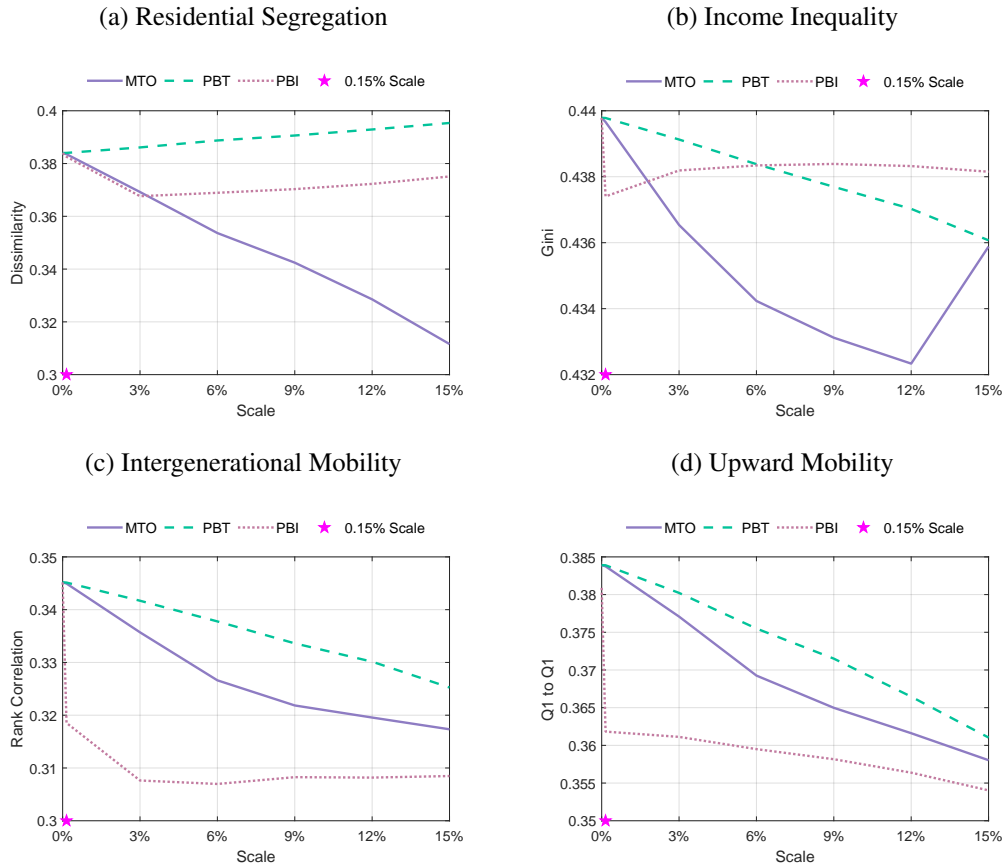
Panels (b), (c), and (d) report the aggregate effects on income inequality, residential segregation, and intergenerational mobility. The case of a fixed eligible pool achieves slightly lower improvements in income inequality and intergenerational mobility, and very similar changes in dissimilarity relative to our baseline case. Overall, the two specifications deliver similar aggregate outcomes.

D.3 More Policies' Comparison

Figure D3 compares the aggregate effects of the MTO, PBT, and PBI policies in their first period of introduction on residential segregation in panel (a), income inequality in panel (b), intergenerational mobility in panel (c), and upward mobility in panel (d). The figure shows that upon impact, the PBI policy reduces inequality and residential segregation more effectively than PBT but less effectively than MTO, while it is

the most effective policy for increasing intergenerational mobility. As we discussed in the main text, turning to dynamic implications uncover larger gains from the PBI policy.

Figure D3: Aggregate effects



D.4 Two Neighborhoods

In the main text we considered a version of our model with three neighborhoods. In this section, we illustrate how the results would change if we considered an alternative version with only two neighborhoods.

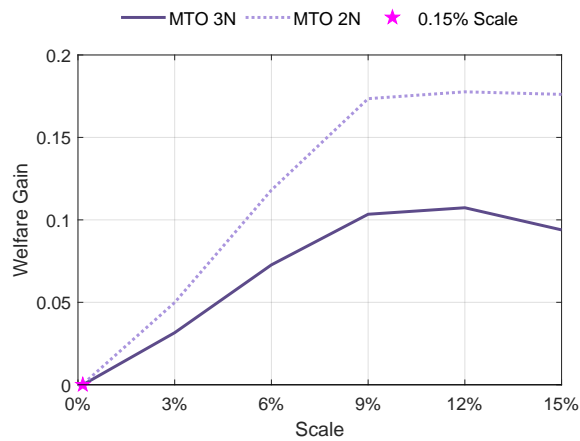
To do this, we consider the following two-neighborhood configuration that changes the definition and size of neighborhood A in the steady state of our model in 1980. Specifically, the population of neighborhood B is split halfway between A and C in the steady state, so we have two neighborhoods, A and C, with respective sizes of 0.3438 and 0.6562 in 1980.⁵¹

⁵¹The results are consistent when considering alternative two-neighborhood specifications, such as (i) a “Small A” case, where the population of neighborhood B becomes part of C in the steady state, so we have two neighborhoods, A and C, with respective sizes of 0.1935 and 0.8065 in 1980; (ii) a “Large A” case, where the population of neighborhood B becomes part of A in the steady state, so we have two neighborhoods, A and C, with respective sizes of 0.4941 and 0.5060 in 1980.

We keep all parameters fixed to our baseline calibration whenever possible, except for the housing supply shifters λ_A and λ_C , which are allowed to change in order to obtain the target size of the two neighborhoods in 1980.

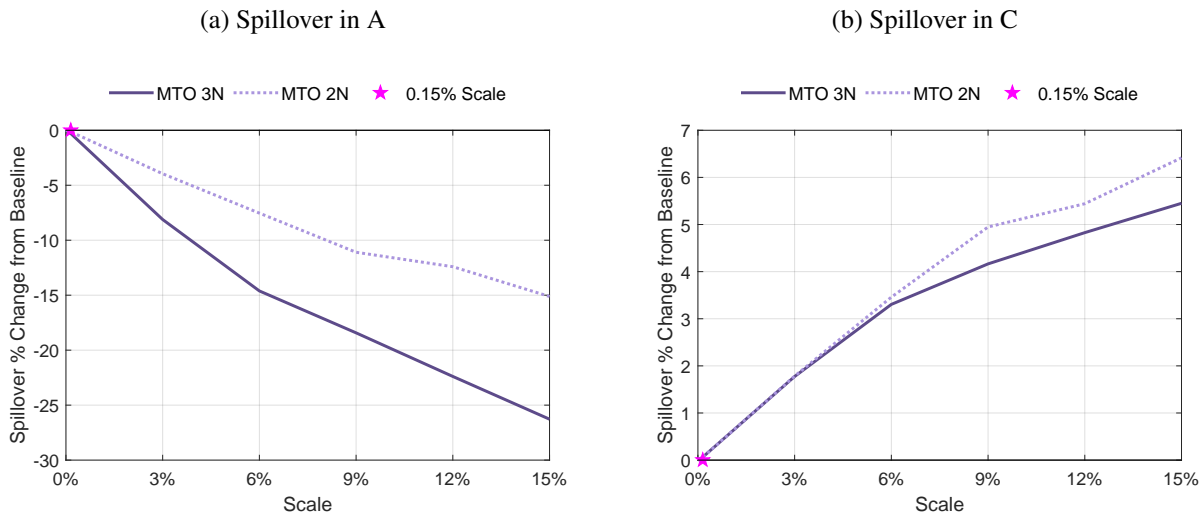
Figure D4 shows the welfare gains from the MTO policy in our baseline case of three neighborhoods and with two neighborhoods. We observe that the model with two neighborhoods delivers higher welfare gains than the three-neighborhoods case.

Figure D4: MTO Welfare Gain 2 vs. 3 Neighborhoods



To understand the welfare effects, in Figure D5 we report the change in the spillover in neighborhood A in panel (a) and C in panel (b) relative to the baseline.

Figure D5: MTO Spillovers 2 vs. 3 Neighborhoods



A key mechanism that emerges is that the lack of the intermediate neighborhood B changes the response of non-recipients to the policy. On the one hand, the lack of an intermediate neighborhood limits the outflow of families from A, despite the decreasing spillover in A as the policy scale increases. This dampens the decline in the spillover of neighborhood A in response to the recipient families moving in. On the other hand, families that move out of A in the presence of the policy need to move to C. Since families arriving from A have higher income than incumbent parents in C on average, they lead to a larger improvement of the spillover in C than in the case of three neighborhoods.

This mechanism results in a lower decline of the spillover in A compared to the three-neighborhoods case, paired with a larger increase in the spillover in C. As a consequence, the welfare gains from the policy are higher than with three neighborhoods.

D.5 Sensitivity to the Altruism Parameter

In the baseline model, parental preferences assign equal weight to own consumption and children’s future income:

$$\log(\theta_{n_t} c_t) + \log(w_{t+1}) + \sigma_\zeta \zeta_{n_t} - \mu I_{n_t \neq b_t}.$$

In this section, we explore the sensitivity of our results to the degree of parental altruism by introducing a parameter φ that governs the relative weight on children’s future income:

$$\log(\theta_{n_t} c_t) + \varphi \log(w_{t+1}) + \sigma_\zeta \zeta_{n_t} - \mu I_{n_t \neq b_t}.$$

When $\varphi < 1$, parents care relatively more about their own consumption; when $\varphi > 1$, they place greater weight on their children’s future income. The baseline model corresponds to $\varphi = 1$. We consider two alternative values: $\varphi = 0.5$ and $\varphi = 1.5$.

The altruism parameter affects both parental decisions—residential choice and educational investment—and the welfare criterion, since welfare is computed using the parents’ utility function. A higher φ increases the marginal return to investing in children’s education and amplifies the welfare consequences of policies that affect children’s future income.

Figure D6 shows the dynamic welfare effects of the three policies at the 3% scale under the two alternative values of φ . When $\varphi = 0.5$ (panel a), the ranking of the three policies is the same as in the baseline: the PBI policy generates the largest long-run welfare gains, followed by PBT, and then MTO. However, the magnitudes are compressed relative to the baseline, as parents place less weight on the children’s income gains generated by the policies. In particular, the MTO policy generates small but positive welfare gains throughout, since the negative effects on non-recipients’ children’s income receive less weight. When $\varphi = 1.5$ (panel b), the differences across policies are amplified. The MTO policy generates growing welfare losses, reaching about -10% in consumption equivalents, as parents place greater weight on the decline in children’s income caused by the deterioration of the spillover in neighborhood A. The PBI and PBT policies

both generate larger welfare gains than in the baseline, with PBI reaching about 18% and PBT about 12% in the long run, as parents value the children’s income gains from these policies more heavily.

Figure D6: Welfare Dynamics: Sensitivity to Altruism Parameter


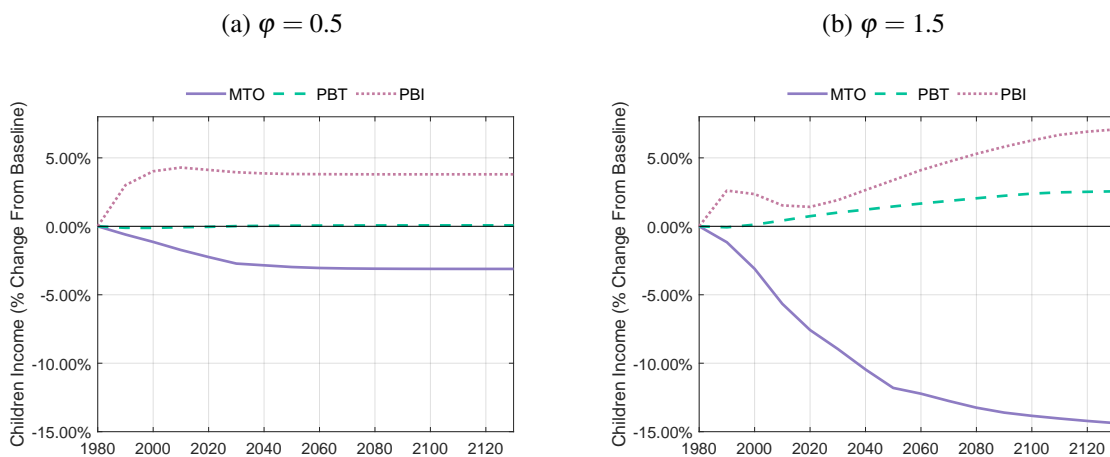


Figure D7 shows the corresponding effects on city-wide average children’s income. When $\phi = 0.5$ (panel a), the PBI policy generates the largest children’s income gains (around 4%), while the MTO policy leads to a decline of about 3% in the long run and the PBT policy has negligible effects on children’s income. When $\phi = 1.5$ (panel b), the MTO policy generates a steeper decline in children’s income, reaching about -15% , reflecting the stronger educational investment response to the deteriorating spillover in neighborhood A. The PBI and PBT policies generate growing income gains, as parents invest more heavily in education in response to the higher returns generated by the policies.

Figure D7: Children’s Income Dynamics: Sensitivity to Altruism Parameter



Overall, the sensitivity analysis confirms that the ranking of the three policies is robust to the degree of parental altruism. The PBI policy dominates in terms of long-run welfare and children’s income gains

under all values of φ . However, the magnitudes of the welfare effects change substantially: higher altruism amplifies both the benefits of the PBI and PBT policies and the costs of the MTO policy.

D.6 Sensitivity to PBI Returns and Durability

The calibration of the PBI parameter ζ in the baseline model relies on estimates from Biasi et al. (2025) on the returns to capital investment in public schools.

Their baseline estimate is that a cumulative \$1,000 increase in raw annual per-pupil spending in low-SES districts raises test scores by 0.0276 standard deviations. They also offer a conversion to the annual “cost” of capital services, which accounts for the durability of the capital investment. To convert from raw spending to capital services, Biasi et al. (2025) account for the fact that capital investments are durable and generate a declining stream of services over time. Assuming 9% annual depreciation over 30 years and assuming that students are exposed to the capital services for five years, they obtain a scaling factor of approximately 2.8. Applying this factor yields a flow value of capital services equal to $0.0276 \times 2.8 = 0.08$ standard deviations per \$1,000 of cumulative spending, which is the estimate we use in our baseline calibration. We choose to use the estimate with this adjustment in our baseline calibration given that our baseline model assumes full depreciation of the investment in every period.

In this section, we consider two robustness exercises that relax these assumptions. First, we recalibrate the PBI parameter ζ using the more conservative raw-spending estimate of 0.0276 standard deviations per \$1,000, rather than the adjusted estimate of 0.08. This yields a calibrated value of $\zeta = 0.28$, and we re-run the dynamic exercise with this parameter while maintaining the assumption of full depreciation. Second, we extend the model to allow for accumulation of the PBI capital stock over time, which is the natural counterpart of using the raw-spending estimate: if the investment is durable, the flow of capital services that it generates should be reflected in the model through a persistent capital stock rather than through a one-time boost to the spillover.

Specifically, we assume that investment creates a capital stock K_{nt} that accumulates over time with geometric depreciation:

$$K_{nt} = (1 - \delta)K_{n,t-1} + I_{nt},$$

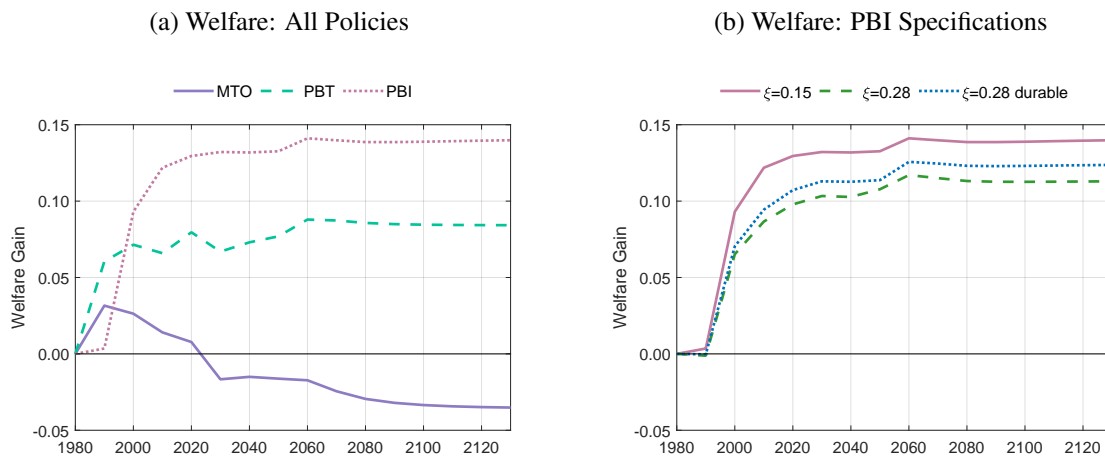
where $\delta \in (0, 1)$ is the per-period depreciation rate and $K_{n0} = 0$. The effective spillover then becomes

$$S_{nt}^* = S_{nt} + \psi K_{nt}^\zeta,$$

so that the spillover depends on the accumulated capital stock rather than on the current investment flow alone. We set the annual depreciation rate to 9%, consistent with Biasi et al. (2025), which over a 10-year model period implies $\delta = 1 - 0.91^{10} \approx 0.61$. Under this specification, approximately 39% of each period’s investment survives to the next period, and the capital stock builds up over time as the government continues to invest in each period.

Panel (a) in Figure D8 reproduces the baseline comparison of the welfare dynamics across the three policies (MTO, PBT, and PBI), showing that the PBI policy generates the largest long-run welfare gains. Panel (b) compares the welfare dynamics of the PBI policy under three alternative specifications: the baseline with $\zeta = 0.15$, the PBI with $\zeta = 0.28$ and full depreciation, and the PBI with $\zeta = 0.28$ and durable capital. A higher ζ dampens the welfare gains from PBI, but the durable capital specification increases these gains in the long run, as the accumulated stock of investment reinforces the spillover effects over time. Importantly, even under the most conservative specification ($\zeta = 0.28$ with full depreciation), the PBI policy dominates MTO and PBT in the long run. The alternative specifications with higher returns and durable capital only strengthen this conclusion, confirming that the PBI policy is the most effective at generating sustained welfare improvements.

Figure D8: Welfare Dynamics: Durable Place-Based Investment

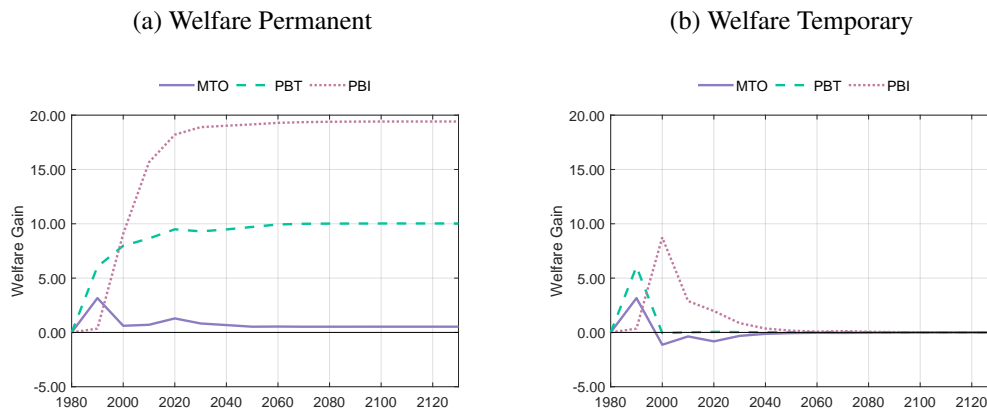


D.7 Temporary Policy

Up to this point, our analysis has focused on permanent neighborhood-specific policies. In this section, we consider the other extreme case, in which such policies are implemented for just one period, in 1990, while maintaining our focus on the scale of 3%.

Figure D9 contrasts the welfare gains of the temporary implementation of the three policies (panel b) with the corresponding permanent versions (panel a), which for convenience replicate the ones in Figure 12. Overall, the results show that although the welfare effects of the temporary policies fade over time, they are most persistent for the PBI policy. In particular, the figure shows that all three policies, when temporary, generate short-lived welfare gains, after which the economies converge back to the benchmark steady state in the long run. The welfare effects of the MTO and PBT policies peak in the year of implementation, whereas those of the PBI policy peak in the subsequent period. This delayed effect arises because the PBI policy operates mostly through children's future income, which materializes with a one-period lag even after the policy has been withdrawn. In addition, the PBI policy exhibits the most persistent effects, lasting up to

Figure D9: Aggregate Effects: Permanent versus Temporary Policy



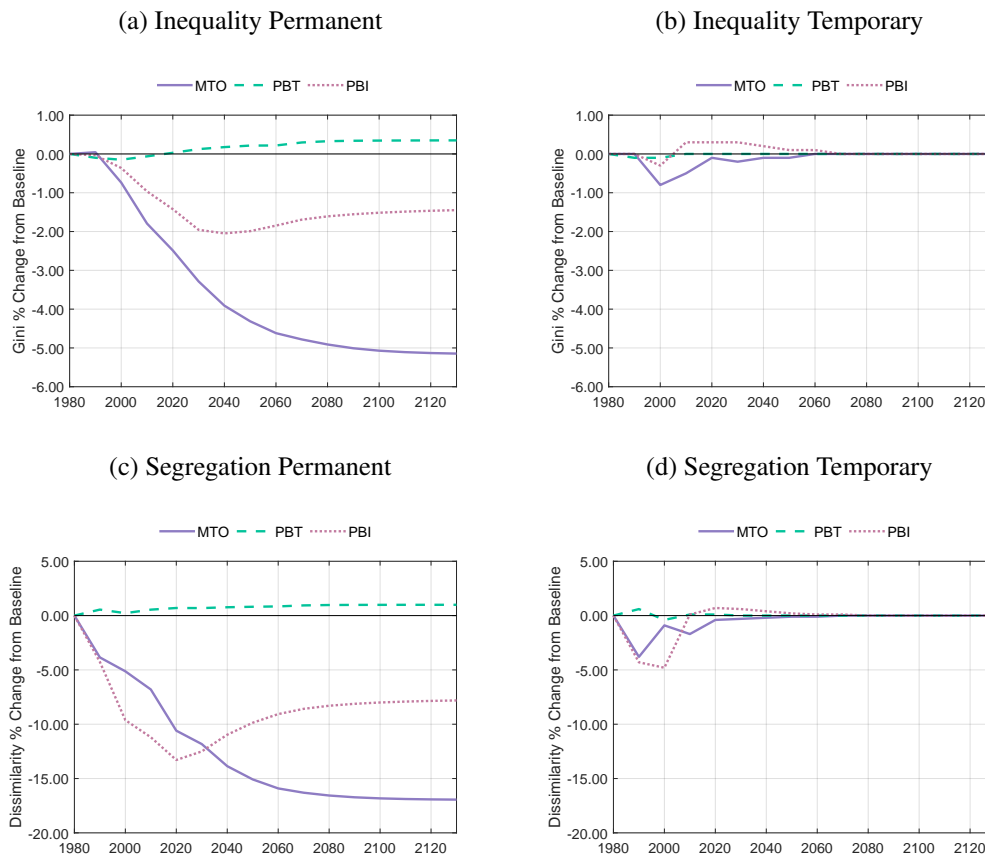
five periods after the policy withdrawal, whereas the benefits of the MTO and PBT policies disappear in the period immediately following the policy phase-out.

Figure D10 compares the effects of the temporary implementations of the three policies on inequality and segregation with those of their permanent counterparts. The figure shows that the effects of temporary policies on inequality and segregation also fade over time. Moreover, under the temporary PBI policy, segregation and inequality initially decrease relative to the baseline, but then they increase, before converging to the initial steady state. This effect happens because the PBI policy allows many high-ability children to achieve higher income than they would in the baseline. Once the policy is withdrawn, the children who benefited from the policy turn into high-income adults who can afford to choose neighborhood A, while neighborhoods B and C are less attractive without the PBI funding in place. This flow of high-income individuals to A increases the size and the spillover of A, inducing a temporary increase in income inequality and segregation.

References

- Agostinelli, Francesco, Margaux Lufade, Paolo Martellini et al., “On the spatial determinants of educational access,” Technical Report 2021.
- , Matthias Doepke, Giuseppe Sorrenti, and Fabrizio Zilibotti, “It takes a village: the economics of parenting with neighborhood and peer effects,” 2020.
- Aliprantis, Dionissi and Daniel R Carroll, “Neighborhood dynamics and the distribution of opportunity,” *Quantitative Economics*, 2018, 9 (1), 247–303.
- Baum-Snow, Nathaniel and Lu Han, “The microgeography of housing supply,” *Journal of Political Economy*, 2024, 132(6), 1897–1946.

Figure D10: Aggregate Effects: Permanent versus Temporary Policy



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.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F Katz, and Christopher Palmer, “Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice,” Technical Report, National Bureau of Economic Research 2019.

—, —, —, —, —, —, and —, “Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice,” *The American Economic Review*, forthcoming.

Biasi, Barbara, Julien Lafortune, and David Schönholzer, “What Works and for Whom? Effectiveness and Efficiency of School Capital Investments Across the US,” *The Quarterly Journal of Economics*, 2025,

140(3), 2329–2379.

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

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

—, **John N. Friedman, and Jonah E. Rockoff**, “Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood,” *American Economic Review*, September 2014, 104 (9), 2633–79.

—, **Nathaniel Hendren, 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.

Chyn, Eric and Diego Daruich, “An Equilibrium Analysis of the Effects of Neighborhood-based Interventions on Children,” *American Economic Review*, 2022, 115(12), 4476–4522.

Clampet-Lundquist, Susan and Douglas S Massey, “Neighborhood effects on economic self-sufficiency: A reconsideration of the Moving to Opportunity experiment,” *American Journal of Sociology*, 2008, 114 (1), 107–143.

Cunha, Flavio, James J Heckman, and Susanne M Schennach, “Estimating the technology of cognitive and noncognitive skill formation,” *Econometrica*, 2010, 78 (3), 883–931.

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

Derenoncourt, Ellora, “Can you move to opportunity? Evidence from the Great Migration,” *American Economic Review*, 2022, 112 (2), 369–408.

Diamond, Rebecca and Tim McQuade, “Who wants affordable housing in their backyard? An equilibrium analysis of low-income property development,” *Journal of Political Economy*, 2019, 127 (3), 1063–1117.

—, —, **and Franklin Qian**, “Who pays for rent control? Heterogeneous landlord response to San Francisco’s rent control expansion,” 2019, 109, 377–380.

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.

Feins, Judith D, Mary Joel Holin, and Anthony A Phipps, *MTO: Moving to Opportunity for Fair Housing*

- Demonstration: Programs Operations Manual*, Abt Associates, 1996.
- Fernandez, Raquel and Richard Rogerson**, “Income distribution, communities, and the quality of public education,” *The Quarterly Journal of Economics*, 1996, *111* (1), 135–164.
- and —, “Public education and income distribution: A dynamic quantitative evaluation of education-finance reform,” *American Economic Review*, 1998, pp. 813–833.
- Fogli, Alessandra, Veronica Guerrieri, and Marta Prato**, “Investigating the American Dream: The Role of Neighborhoods,” *mimeo*, 2026.
- , —, **Mark Ponder, and Marta Prato**, “The End of the American Dream? Inequality and Segregation in US Cities,” *Journal of Political Economy*, 2026, *134*(4), 1110–1158.
- , —, **Martin Garcia-Vazquez, and Marta Prato**, “Neighborhood Segregation and Endogenous Racial Bias,” *mimeo*, 2026.
- Galiani, Sebastian, Alvin Murphy, and Juan Pantano**, “Estimating neighborhood choice models: Lessons from a housing assistance experiment,” *American Economic Review*, 2015, *105* (11), 3385–3415.
- Gaubert, Cecile, Patrick M Kline, and Danny Yagan**, “Place-based redistribution,” Technical Report, National Bureau of Economic Research 2021.
- Goering, John, Joan Kraft, Judith Feins, Debra McInnis, Mary Joel Holin, and Huda Elhassan**, “Moving to Opportunity for fair housing demonstration program: Current status and initial findings,” *Washington, DC: US Department of Housing and Urban Development*, 1999.
- Goldin, Claudia Dale and Lawrence F Katz**, *The race between education and technology*, Harvard University Press, 2009.
- Imberman, Scott A, Adriana D Kugler, and Bruce I Sacerdote**, “Katrina’s children: Evidence on the structure of peer effects from hurricane evacuees,” *American Economic Review*, 2012, *102* (5), 2048–82.
- Katz, Lawrence F, Jeffrey R Kling, and Jeffrey B Liebman**, “Moving to opportunity in Boston: Early results of a randomized mobility experiment,” *The Quarterly Journal of Economics*, 2001, *116* (2), 607–654.
- Kling, Jeffrey R, Jeffrey B Liebman, and Lawrence F Katz**, “Experimental analysis of neighborhood effects,” *Econometrica*, 2007, *75* (1), 83–119.
- Lavy, Victor, M Daniele Paserman, and Analia Schlosser**, “Inside the black box of ability peer effects: Evidence from variation in the proportion of low achievers in the classroom,” *The Economic Journal*, 2012, *122* (559), 208–237.
- Ludwig, Jens, Greg J Duncan, Lisa A Gennetian, Lawrence F Katz, Ronald C Kessler, Jeffrey R Kling, and Lisa Sanbonmatsu**, “Long-term neighborhood effects on low-income families: Evidence from Moving to Opportunity,” *American economic review*, 2013, *103* (3), 226–231.
- Redding, Stephen J and Daniel M Sturm**, “Estimating neighborhood effects: evidence from war-time

destruction in London,” Technical Report 2016.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Raymond Owens III, “Housing externalities,” *Journal of political Economy*, 2010, 118 (3), 485–535.

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

—, “Peer effects in education: How might they work, how big are they and how much do we know thus far?,” in “Handbook of the Economics of Education,” Vol. 3, Elsevier, 2011, pp. 249–277.

Saiz, Albert, “The Geographic Determinants of Housing Supply,” *The Quarterly Journal of Economics*, 08 2010, 125 (3), 1253–1296.

Sanbonmatsu, Lisa, Lawrence F Katz, Jens Ludwig, Lisa A Gennetian, Greg J Duncan, Ronald C Kessler, Emma K Adam, Thomas McDade, and Stacy T Lindau, “Moving to opportunity for fair housing demonstration program: Final impacts evaluation,” 2011.