

An Equilibrium Analysis of the Effects of Neighborhood-based Interventions on Children*

Eric Chyn and Diego Daruich

October 17, 2022

Abstract

This paper studies housing vouchers and urban redevelopment programs by incorporating neighborhood effects into a general equilibrium overlapping-generations model with endogenous location choice and child development. We calibrate the model using U.S. data and show that simulated predictions match reduced form evidence from the literature. Our analysis finds that large-scale implementations of vouchers and place-based subsidies both result in long-run welfare gains by reducing inequality and generating skill improvements that offset higher taxation and other GE effects. Although vouchers lead to larger welfare gains on average, we find that place-based subsidies may be preferable in cities with constrained housing supply.

Keywords: Neighborhood effects, housing vouchers, intergenerational mobility.

JEL Codes: J13, R13, R23, I31.

***Chyn:** Department of Economics, University of Texas at Austin, and the National Bureau of Economic Research, Email: eric.chyn@austin.utexas.edu. **Daruich:** Marshall School of Business, University of Southern California, Email: daruich@usc.edu. We are grateful for helpful comments and suggestions that we received from Treb Allen, Selahattin Imrohorglu, Jesse Gregory, Claudia Olivetti, Alessandra Fogli, Jonathan Heathcote, and seminar participants at USC Marshall, the Junior Micro Macro Labor Conference at Columbia University, the CAF Development Bank of Latin America, Virginia Commonwealth University, the Stanford Institute for Theoretical Economics, William and Mary University, UC-Riverside, Universidad de San Andres, the Inter-American Development Bank, the Opportunity and Inclusive Growth Institute (at the Federal Reserve Bank of Minneapolis), the Society for Economic Dynamics, the Western Economic Association International, and the NBER Summer Institute 2022.

1 Introduction

An emerging literature demonstrates that neighborhoods have important impacts on long-run outcomes of children. Recent analysis of the Moving to Opportunity (MTO) experiment finds substantial improvement in earnings and other outcomes for young children whose families moved to low-poverty neighborhoods using subsidized housing vouchers (Chetty et al., 2016). Similarly, Chyn (2018) finds notable long-run gains for children whose families were forced to relocate to less disadvantaged areas due to public housing demolitions.

In light of this evidence, a natural question is: How should governments design policies to improve neighborhood quality for children? The answer to this question depends on general equilibrium (GE) responses that are not well-captured by highly credible but short-run and relatively small experimental studies. For example, the benefits of encouraging poor families with children to move to low-poverty areas may be diminished if the characteristics of more advantaged neighborhoods change in response over time.¹

This paper provides a new assessment of the equilibrium effects of housing mobility programs and government policies that aim to revitalize disadvantaged neighborhoods. We study a spatial equilibrium model that features overlapping generations and incorporates endogenous childhood development. Our model extends on seminal work that theoretically studies inequality and neighborhoods (Benabou, 1996*a,b*; Durlauf, 1996; Fernandez and Rogerson, 1996, 1998).

Our framework consists of three main building blocks that are key for a comprehensive assessment of the effects of policies that shape childhood exposure to disadvantaged neighborhoods. The first is that parental choices are important for child outcomes. An individual's productivity depends on skills that are influenced by parental choices. Specifically, parents choose one of two neighborhoods and make time investments into their children. Neighborhood quality matters due to local externalities. The second building block is a GE life-cycle Aiyagari framework that features endogenous labor supply and embeds parental investments and inter-generational linkages. Wage shocks in this block of the model increase income inequality and help explain why parents may be unable to move to a more advantaged neighborhood. The GE forces also allow us to account for the effects of public policies on prices in the economy—i.e., housing costs, capital returns, and wages. Finally, the third building block is the government which funds policy interventions by levying income taxes. Taxes have distortionary effects due to the endogeneity of labor supply and human capital in our model.

We estimate the model using simulated method of moments to match recent data on the

¹Prior research has justified this concern by showing that changes in neighborhood demographics may cause out-migration of incumbents and alter the distribution of public goods (Boustan, 2010; Derenoncourt, 2019).

geography of opportunity in the U.S. We map the neighborhoods in our model to Census tracts in the U.S. For each commuting zone (CZ), we divide Census tracts into two groups according to income per capita, the bottom 10 percent (to capture the usual targets of anti-poverty housing and place-based policies) and top 90 percent, and then average across CZs. In addition to matching standard moments (e.g., the average hours worked), we target those that are informative about income and parental investment by neighborhood as well as moments related to neighborhood externalities. For the latter, we rely on data on long-run outcomes and childhood neighborhoods from the Opportunity Atlas (Chetty et al., 2018). Our model requires us to specify explicitly how time and neighborhood characteristics aggregate to form “parental investments.” We do this via a CES aggregator and estimate the parameters of this function by matching the income of children who grow in different neighborhoods, the average amount of quality time parents spend with their children, and the differences in time across income groups.

As validation exercises, we show that simulated predictions from the calibrated model match reduced form evidence from experimental and quasi-experimental studies. First, we demonstrate that the model is in line with experimental estimates of the impact of moving from Chetty et al. (2016). Chetty et al. studied disadvantaged families that received housing vouchers that could only be used in a low-poverty neighborhood through the MTO randomized control trial (RCT). We study an equivalent program within our model that mimics the features of the small-scale and short-run nature of the RCT.² We find that the model-generated impacts on the labor market outcomes of children treated by the intervention are similar to those from MTO. Second, we find that a simulated version of a place-based wage subsidy program generates impacts that match evidence from Busso et al. (2013). They use quasi-experimental methods to study the Empowerment Zone (EZ) program, a federal policy in the U.S. that provided incentives (e.g., tax credits for employing local workers) to encourage development of disadvantaged urban and rural communities. We simulate the EZ program as a place-based wage subsidy program and obtain predicted earnings gains for adults that match Busso et al. (2013). The results from these two exercises provide evidence that the model is in line with the most credible reduced-form evidence on the impacts of housing vouchers and place-based incentive programs.

After this validation exercise, we begin by studying the long-run effects of housing voucher programs taking into account GE effects and financing from progressive labor taxes. We explore versions of the housing voucher policy which differ in terms of three characteristics: (1) the voucher subsidy rate; (2) an eligibility restriction in terms of the individuals hourly wage; and (3) an eligibility restriction based on the presence of children (which is based on age given the

²That is, we simulate effects for a single generation while holding prices and neighborhood qualities fixed.

model's exogenous fertility assumption). The highest steady-state welfare gains are achieved with a policy that has a full subsidy rate and targets households that have children and wages below the ninth decile (i.e., the 90th percentile). This voucher program generates a 4.3 percent increase in consumption equivalent units, despite the fact that the average marginal tax rate must increase by 16 percent to fund the voucher program. As an additional 12.9 percent of children move to the better neighborhood, labor productivity increases by 1.4 percent. In addition, we find that the voucher program has consequences for inequality and upward mobility. Specifically, the program leads to a reduction in the variance of log-after-tax-lifetime-earnings of 2.7 percent along with an increase in upward mobility by 29.4 percent.³ The inequality effect is roughly as large as half of the difference in the variance of log-income between Sweden and the U.S. The effect on upward mobility is approximately equal to one-half of its standard deviation across US Census tracts.

We decompose these results—particularly the welfare gains of 4.3 percent—into five key equilibrium effects. We begin with a simulation of a short-run partial-equilibrium version (i.e., implemented on a small group for a single generation) of our highest steady-state welfare voucher program to mimic an RCT. This scenario leads to welfare gains of 5.7 percent. An additional simulation that accounts for tax increases needed to finance the voucher program shows that welfare gains in a large-scale intervention decrease by 1.3 percentage points, i.e., approximately 25 percent. The rent and neighborhood quality equilibrium effects (e.g., lower-income individuals tend to move to the advantaged neighborhood) jointly reduce the benefits by 3.3 percentage points, half of the simulated RCT gains. Long-run intergenerational dynamics and equilibrium effects on wages and the interest rate almost perfectly compensate for this, increasing welfare gains by 3.2 percentage points. Long-run dynamics increase gains because investing in a child not only improves their skills but also creates better parents for the next generation. In sum, the five equilibrium effects ultimately make the long-run general-equilibrium welfare gains about 25 percent smaller than those from the short-run partial-equilibrium version of the program.

Next, we examine long-run GE effects of place-based policies. As in our validation exercise, we study a neighborhood-specific wage subsidy program and explore versions of the program that vary the level of the subsidy. The highest steady-state welfare gains are achieved with a 11 percent wage subsidy. This policy achieves a 0.5 percent increase in consumption equivalence terms, notably smaller benefits than what is possible with a voucher program. As a result of the subsidy, there is substantial resorting to the disadvantaged neighborhood, and the share of children living in the advantaged area decreases by 7.4 percent. Income inequality decreases by

³Note that we measure upward mobility as the probability that a child reaches the top 20 percent of the income distribution given that had parents with income at the bottom 20 percent of the distribution.

8.7 percent and upward mobility increases by 21.6 percent.

Our decomposition analysis for the wage subsidy program shows that most of the equilibrium forces that we consider have relatively important roles in determining the 0.5 percent welfare gain. When the program is implemented in a short-run (so the subsidy is not provided to future children) partial equilibrium world, the impact on welfare (calculated only for children) is negative at 1.4 percent since the wage subsidy induces parents to relocate to the disadvantaged neighborhood. When the equilibrium allows for neighborhood quality and rental price adjustments, welfare losses are eliminated. The tax increases needed to finance the program reduce welfare gains by 0.2 percentage points. Implementing the program for the long-run further increases welfare by 0.7 percentage points because neighborhood quality increases more in the long run, and children in this scenario take advantage of the wage subsidy when they reach adulthood.

To assess the sensitivity of these findings regarding vouchers and place-based subsidies, we follow Andrews et al. (2017) and Elenev et al. (2020) and examine the robustness of our welfare results to changes in the parameters used in our policy simulations. This analysis reveals two main findings. First, empirically reasonable individual changes to most parameters do not substantially affect the welfare gains achieved under either policy. Second, notably large changes in the value of the housing supply elasticity can potentially reshape the debate over the merits of vouchers and place-based subsidies. For our main analysis, we rely on a housing supply elasticity estimate that is representative of the *average* major U.S. metropolitan area. A natural alternative is to consider the case when housing supply is highly constrained as in locations such as San Francisco or New York. The results from our sensitivity analysis suggest that reducing the supply elasticity to match the level observed in the most land-constrained cities may tip the balance so that place-based subsidies generate higher long-run welfare gains relative to vouchers.

Why do government policies that shape exposure to high-quality neighborhoods increase welfare? There are two main explanations within our model. First, neighborhood externalities create a role for place-based policies because the government accounts for the fact that individual work choices affect skills of children. Location-based wage subsidies are a means of increasing this positive externality. Second, the main channel for welfare improvement through housing vouchers lies in the government's capacity to make up for the absence of intergenerational borrowing—i.e., a parent's inability to borrow against their child's future income. For example, a poor parent who could invest in their child's development by moving would want to smooth consumption intergenerationally. The inability to make this type of transfer reduces the incentive to move. Housing vouchers can be thought of as using taxation to address this market failure. In sum, these two factors imply that the government can use housing and urban development poli-

cies to invest in children and tax them later once they become adults.

In terms of distributional impacts, a natural consideration is that the programs we consider may have heterogeneous effects on welfare for the adults alive at the introduction of the policy. We analyze transition dynamics and find both policies have relatively concentrated gains and a majority of (incumbent) adults would vote against them. For housing vouchers, the gains are concentrated among young cohorts. In contrast, the wage subsidy program generates gains mainly for those living in the low-income neighborhood. Overall, the lack of support for both policies suggests that interventions that generate long-run gains by improving neighborhood quality for children may have important political economy tradeoffs.⁴

Our analysis and findings contribute to a large and growing literature that studies neighborhoods and government policies to promote urban development. A number of recent studies have focused on providing credible reduced form evidence on the effects of moving using housing vouchers (Kling et al., 2007; Chetty et al., 2016; Chyn, 2018) and the neighborhood-level impacts of place-based policies such as Empowerment Zones (Busso et al., 2013). Relatively few studies use equilibrium frameworks to study housing assistance policies. Diamond and McQuade (2019) and Davis et al. (2019) study the effects of government programs that construct low-income housing. Closer to the concerns of this paper, Davis et al. (2021) study rental vouchers and equilibrium sorting behavior.⁵ Our analysis complements these prior works by studying housing mobility and place-based policies in a single framework that accounts for labor supply responses and taxation—two features that we find are important for understanding large-scale equilibrium responses.

Finally, this paper is also closely related to an emerging literature in macroeconomics that quantitatively studies location choice, inequality, and children. Important work by Fogli and Guerrieri (2019), Zheng and Graham (2020) and Eckert and Kleineberg (2021) similarly use spatial equilibrium models to study child development but differ in at least two ways.⁶ First, we focus on

⁴As in other studies that find gains are driven by benefits for future cohorts (e.g., Daruich, 2020), one may expect that these policies could be acceptable to a majority of individuals if the government borrows to obtain temporary financing (while increasing taxation in the future to pay for the government debt).

⁵A key exercise in Davis et al. (2021) studies the optimal design of housing vouchers and finds that maximizing the impact of vouchers requires restricting vouchers to be redeemable only in high opportunity neighborhoods. Our model features only two neighborhoods so we are unable to consider how location restrictions on vouchers may impact welfare gains. That said, simulation results show that our simplified model is able to replicate experimental estimates of the effects of voucher moves through the MTO program. We take this validation result as suggesting that the type of voucher-based moves that we study are relevant for understanding voucher policy.

⁶Agostinelli et al. (2020) and Aliprantis and Carroll (2018) also use equilibrium models that feature child development to study neighborhoods and relocation policies, but both works differ from our analysis in their focus. The former studies peer selection and short-run effects whereas we examine long-run dynamics, taxation, and housing equilibrium effects. The latter studies the consequences of allowing for mobility and equalizing neighborhoods productivity while abstracting from labor supply and taxation effects—two features that we find to be quantitatively important for our analysis of the welfare impacts of housing and neighborhood investment programs.

related but distinct questions on the effects of residential choice. While [Fogli and Guerrieri \(2019\)](#) study the contribution of segregation to increases in U.S. inequality since the 1980s, we use our calibrated model to study counterfactual welfare gains. Similar to our study, [Zheng and Graham \(2020\)](#) and [Eckert and Kleineberg \(2021\)](#) consider welfare questions, but we focus on different policies. Their analysis centers on the effects of equalizing school funding whereas we study the effects of housing vouchers and place-based incentive programs, two prominent types of government policies in countries around the world.⁷ Second, [Zheng and Graham \(2020\)](#) and [Eckert and Kleineberg \(2021\)](#) use models that have more spatial heterogeneity and allow for greater location choice. The main benefit of their approach is that they can evaluate policy effects on more locations. In contrast, our model allows for a larger range of equilibrium effects—through changes in wages, capital returns, taxes, as well as housing costs and neighborhood qualities—in the long-run steady state and during the transition after policy adoption. A key advantage of this approach is that our analysis accounts for the costs of raising taxes to pay for policies of interest. In addition, by incorporating the analysis of transitional dynamics, we can assess issues of political economy. This feature of our analysis delivers one of our core findings in that the policy with the largest long-run welfare gains may not have the broadest political support.

2 Motivating Facts


To motivate our model and analysis, this section highlights key findings regarding the distribution of economic outcomes across neighborhoods. We focus on neighborhood median household income and long-run outcomes of children as measured by household income. The data on child household income comes from the Opportunity Atlas ([Chetty et al., 2018](#)) which measured income using IRS records on mean earnings in 2014-2015 when a child was between ages 31-37. The data provide measures of mean household income for children who have parents at the 25th and 75th percentile of the national income distribution. All measures are available at the neighborhood level as defined by U.S. Census tracts. Tracts are small geographic units that have an average population of 4,250 persons.

The main pattern that we note is that there is a significant correlation between the economic outcomes of adults in a neighborhood and the long-run outcomes of children growing up in these areas. Panel A of [Figure 1](#) is a binned scatterplot that illustrates this relationship for children with low-income parents after controlling for commuting zone fixed effects to account for broad

⁷In the U.S., approximately 2.3 million households receive assistance in the form of a Section 8 voucher each year ([Collinson et al., 2015](#)). Large-scale housing subsidies are also prevalent in European countries ([Salvi et al., 2016](#)). In terms of place-based policies, the U.S. currently spends nearly \$60 billion annually on such programs ([Bartik, 2020](#)). As noted by [Neumark and Simpson \(2015\)](#), a number of European countries also use place-based policies to aid municipalities that have high rates of unemployment or poverty.

differences across metropolitan areas.⁸ These results show that children from low-income families who grow up in tracts where adults have higher incomes have notably higher incomes later in life. A simple regression shows that every \$1,000 increase in the median household income of adults in an area is associated with an increase in the expected household income of poor children by roughly \$3,300. Panel B shows that there is a similar positive correlation for children with high-income parents.

Figure 1: Correlations Between Neighborhood Income and Expected Income of Children



Notes: Panel (a) is a binned scatterplot of the relationship between neighborhood median household income from the 1990 U.S. Decennial Census (x -axis) and estimates of mean household income for children who grew up in the tract and had parents with household income (y -axis) for children with parents at the 25th percentile of the national income distribution. Panel (b) provides similar statistics for children with parents at the 75th percentile of the national income distribution. Statistics are based on data from the Opportunity Atlas (Chetty et al., 2018). All measures are specific to children who were born in the 1978-83 cohorts. The underlying data has information on the mean household income ranks for children. We use the national income distribution statistics to convert income ranks into 2015 U.S. dollars. All binned scatterplot results control for commuting zone fixed effects.

Recent studies provide compelling evidence that the correlation between neighborhood economic conditions and the long-run economic outcomes of children is largely driven by causal effects. Chetty et al. (2016) and Chyn (2018) find that moving out of high-poverty neighborhoods has large positive benefits for children living in severely distressed public housing projects. Chetty and Hendren (2018a) use tax records for 7 million families that move across commuting zones (CZs) and find notable benefits from relocating to more advantaged areas. Their analysis suggests that a young child who moved at birth to a better area and stayed there for 15 years would pick up 60 percent of the difference in permanent resident outcomes between their origin and destinations.

⁸Appendix Figure A1 provides two maps of Chicago to provide another illustration of the correlation between neighborhood poverty and long-run child outcomes.


3 Model

The model has three main components. First, the long-run outcomes of children depend on parental choices. An individual's earnings depend on skills that are determined in childhood. Parents affect the skills of their children by choosing neighborhoods and time investments. Neighborhoods matter for children due to endogenous local spillovers. Specifically, local spillovers occur because the skills of children increase when they grow up in an area with higher income per capita. Second, the economy is modeled using a GE life-cycle Aiyagari framework featuring wage uncertainty and incomplete markets. The model features distortive taxation by allowing for endogenous labor supply. A representative firm combines capital and labor from workers who vary by skill to produce a final consumption good. Third, the government levies taxes on consumption, labor, and capital to finance lump-sum transfers and retirement benefits to individuals.

3.1 Individual Choices and Timing

The model assumes a dynastic framework with 20 age periods and three main stages: childhood, working adulthood, and retirement. Figure 2 shows the life cycle of an individual. Periods are four years long. Let j denote the age in each period (e.g., $j = 1$ refers to ages 0–3). From $j = 1$ to $j = 4$, a child lives with their parents in neighborhood n , and they do not make any choices. In our stylized model, the child reaches adulthood and achieves independence at the beginning of $j = 5$ (age 16). At independence, the individual's state variables include their neighborhood n , savings a (from parental transfers), skills θ , and an idiosyncratic moving cost κ .

Figure 2: Model Timeline: A Dynastic Framework with Three Stages



Notes: This figure illustrates key events and the three main stages of life for an agent in the model.

Each period is divided in two parts. In the first part, the agent chooses a neighborhood. Each neighborhood is associated with a rent τ_n (and, potentially, a moving cost). Having selected a neighborhood, the second part of the period occurs. When individuals are young, they are in the working stage so they choose their savings, consumption, and labor supply (where idiosyncratic,

uninsurable risk makes labor income stochastic). Individuals can borrow up to a limit and save through a non-stage-contingent asset. At $j = 8$, the individual becomes a parent and new decisions must be made. For four periods (i.e., until their child is 16 years old), they decide how much time to invest in child development. Time investments and neighborhood choice both determine the child's skills. Before the child becomes independent, the parent also makes a transfer to the child. Once the agent enters the period when $j = 17$, they enter the retirement stage. At this time, agents have two sources of income: savings and government provided retirement benefits.

3.1.1 Working Stage Decisions

During the working stage, individuals consume c , save a' , and choose labor supply h in the second part of each period. These choices depend on the individual's level of assets a , level of skills θ , current neighborhood location n (which is chosen previously during the first part of each period as detailed below), and a stochastic labor efficiency parameter η . Formally, the value function during the working stage when individuals do not have children is given by:

$$V_j(a, \theta, n, \eta) = \max_{c, a', h} \left\{ u(c, h) - \bar{v}_n + \beta \mathbb{E} \left[\widehat{V}_{j+1}(a', \theta, n, \eta') \right] \right\}, \quad (1)$$

$$c(1 + \tau_c) + \tau_n + a' - (y - T(y)) - \omega = \begin{cases} a(1 + r(1 - \tau_a)) & \text{if } a \geq 0 \\ a(1 + r^-) & \text{if } a < 0 \end{cases}$$

$$y = w_n E_j(\theta, \eta) h, \quad a' \geq \underline{a}_j, \quad 0 \leq h \leq 1, \quad \eta' \sim \Gamma_j(\eta).$$

Individuals receive a flow utility given by the function $u(c, h)$ which depends on consumption and labor supply in addition to receiving utility from their neighborhood due to a fixed (exogenous) amenity value \bar{v}_n . An individual can borrow up to an age-specific limit \underline{a}_j by paying interest at rate r . Individuals can also save for a rate of return r . We assume that individuals cannot commute, implying there are neighborhood specific wages denoted by w_n . As discussed further in Section 3.2, the lack of commuting in our model has no impact on our subsequent analysis since wages in the decentralized equilibrium are equal across neighborhoods due to free mobility of capital. Wages are scaled by the function $E_j(\theta, \eta)$, which is an age-specific function of the individual's skills θ and the idiosyncratic labor efficiency η . Finally, individuals pay linear taxes on consumption (given by τ_c) and capital income (τ_a), pay a non-linear (which are progressive as described below) tax on labor income ($T(y)$), and receive lump-sum government transfer ω .

In the first part of each period, individuals choose where to live taking into their expected utility value (which depends on their current state variables and rent costs, as represented above), and moving costs. During the first period of independence ($j = 5$), we specify that the moving cost is heterogeneous—this assumption will help us capture the fact that younger individuals

are more likely to live in lower-income neighborhoods. Given the neighborhood location n in the period when $j = 4$ (chosen by one's parents), the value function determining agent's first location choice at independence (i.e., age-period $j = 5$) is given by:

$$\widehat{V}_{j=5}(a, \theta, n, \eta, \kappa) = \max_{n' \in \{1,2\}} V_{j=5}(a, \theta, n', \eta) - \kappa 1(n' \neq n),$$

where κ is the stochastic utility cost of moving. As specified in Equation 3, we assume that κ is normally distributed with mean $\bar{\kappa}$ and standard deviation σ_κ . Of course, this cost is only incurred when an individual chooses a new neighborhood (i.e., $n' \neq n$).

From $j = 6$ until retirement (which starts at $j = 17$), the individual's optimization problem in the first part of each period (except for parenthood as described below) is similar to Equation 1. The main difference from the first period of independence is that the location choice involves a fixed moving cost $\bar{\kappa}$. Hence, the value function is given by:

$$\widehat{V}_j(a, \theta, n, \eta) = \max_{n' \in \{1,2\}} V_j(a, \theta, n', \eta) - \bar{\kappa} 1(n' \neq n).$$

Note that, while there are no moving cost shocks, wage shocks η can induce workers to move.

3.1.2 Parental Investment and Child Development

The individual's problem changes when a child is born at the exogenously given fertility age-period $j = 8$ (age 28). We assume that each individual has one child. As in [Barro and Becker \(1989\)](#), parents are altruistic, and they care about the child's with the weight $\tilde{\beta}$. Parents invest in children's skills while they are young ($j = 8 - 11$) and give them an asset transfer once they are about to become independent ($j = 12$).

Children are born with skills θ_k that are potentially correlated with parent's skills. To be in line with the estimates from [Cunha \(2013\)](#), we assume that skills are a vector that includes cognitive $\theta_{c,k}$ and non-cognitive $\theta_{nc,k}$ components. During each period of parenthood ($j = 8 - 11$), the individual chooses the number of hours τ to invest in the child's development of skills.

In addition to time investment, the skill development of children also depends on neighborhood quality. We summarize neighborhood quality as a single index measure s_n . We assume that this spillover effect is determined by per capita total income (the sum of capital and labor) for those living in neighborhood n . Note that we include all residents—those working and not working—in this calculation. Intuitively, this allows our measure to capture the idea that the fraction of children in a neighborhood matters. In addition to ideas related to economic resources per child, previous work in sociology highlights that adults within a neighborhood play a key role in promoting community social organization by supervising children and limiting deviant behavior

(Sampson and Groves, 1989; Sampson et al., 2002). In this way, neighborhoods where adults are a larger fraction of the population may be particularly beneficial to children.

Our focus on income captures a number of standard theoretical mechanisms thought to drive neighborhood effects. Areas with richer parents typically have higher quality schools due to the local financing of public schools (Howell and Miller, 1997; Hoxby, 2001; Biasi, 2019).⁹ In addition, children may benefit from growing up with highly productive adults due to role model effects (Wilson, 1987). More generally, our choice of representing effects in terms of earnings broadly follows prior studies that proxy for neighborhood quality using measures of local area income or poverty rates (Kling et al., 2007; Chetty and Hendren, 2018b).

We model skill development θ_k using two nested constant elasticity of substitution (CES) functions that determine the influence of parent time and neighborhood spillovers. The outer CES is based on Cunha et al. (2010) and allows a child's skills in the next period θ'_k to depend on current skills, parental skills θ , parental investments I , and an idiosyncratic shock v . The inner CES function determines I and explicitly incorporates τ and s_n .

Formally, we assume that the problem of parents in age-periods $j = 8 - 11$ is:

$$\begin{aligned}
V_j(a, \theta, n, \eta, \theta_k) &= \max_{c, a', h, c_k, \tau} u(c, h) - \bar{v}_n + \tilde{\beta}u(c_k, 0) - v(\tau) + \beta\mathbb{E}\left(\widehat{V}_{j+1}(a', \theta, n, \eta', \theta'_k)\right), \quad (2) \\
(c + c_k)(1 + \tau_c) + 2\tau_n + a' - (y - T(y)) - \omega &= \begin{cases} a(1 + r(1 - \tau_a)) & \text{if } a \geq 0 \\ a(1 + r^-) & \text{if } a < 0 \end{cases} \\
y = w_n E_j(\theta, \eta) h, \quad a' \geq \underline{a}_j, \quad 0 \leq h + \tau \leq 1, \quad \eta' \sim \Gamma_j(\eta) \\
\theta'_{q,k} &= \left[\alpha_{1,q,j} \theta_{c,k}^{\rho_{q,j}} + \alpha_{2,q,j} \theta_{nc,k}^{\rho_{q,j}} + \alpha_{3,q,j} \theta_c^{\rho_{q,j}} + \alpha_{4,q,j} \theta_{nc}^{\rho_{q,j}} + \alpha_{5,q,j} I^{\rho_{q,j}} \right]^{1/\rho_{q,j}} e^{v_q} \\
v_q &\sim N(0, \sigma_{q,j,v}), \quad q \in \{c, nc\} \\
I &= \bar{A}_j \left[\alpha_{I,j} f(s_n)^\gamma + (1 - \alpha_{I,j}) \tau^\gamma \right]^{1/\gamma}
\end{aligned}$$

Aside from parental time investment, parents also decide on children's consumption c_k , which we assumed is valued by the same utility function as adults consumption weighted by the altruism parameter $\tilde{\beta}$. Given that we estimate rent costs in a per-person basis in Section 4, we assume that rent doubles when a child is present. Regarding the skill development function, the parameter $\rho_{q,j}$ determines the substitutability of ability inputs in the outer CES function for $q \in \{c, nc\}$. The substitutability of parental time investments and neighborhood quality is determined by the parameter γ in the inner CES function.

⁹Prior evidence suggests that schools are not the only mechanism that can generate neighborhood effects. As noted in Chetty et al. (2018), there is substantial variance in child outcomes across Census tracts within the same school attendance zone, and schools account for less than half of the observed variance across tracts within a county.

As in other periods, the individual can move at the beginning of each period. Differently from previous periods, however, the value function for that choice incorporates the children's skills θ_k :

$$\widehat{V}_j(a, \theta, n, \eta, \theta_k) = \max_{n' \in \{1, 2\}} \{ \mathbb{E}(V_j(a, \theta, n', \eta, \theta_k) - \bar{\kappa} 1(n' \neq n)) \}.$$

3.1.3 Child Independence

Before the parent reaches age-period $j = 12$ (i.e., when they are age 44), they choose a monetary transfer \hat{a} given to the child. We model this as a sub-period taking place before the child becomes independent (at age 16), with a value for the agent defined by V_{Transfer} :

$$V_{\text{Transfer}}(a, \theta, n, \eta, \theta_k) = \max_{\hat{a}} \widehat{V}_{j=12}(a - \hat{a}, \theta, n, \eta) + \tilde{\beta} \mathbb{E}_{\eta_k, \kappa} \left(\widehat{V}_{j'=5}(\hat{a}, \theta_k, n, \eta_k, \kappa) \right), \quad (3)$$

$$\hat{a} \geq 0, \quad \kappa \sim N(\bar{\kappa}, \sigma_\kappa), \quad \eta_k \sim \Gamma_{j'=5}.$$

Importantly, the transfer \hat{a} must be non-negative: the parent cannot leave debt to their child nor borrow against the child's future income. When making this choice, the parent knows their own income shock η . They are not aware of the child's income shock η_k or stochastic moving cost draw κ . Note that, unlike Equation 2, the value function at this stage includes the continuation value of the child $\widehat{V}_{j'=5}$ where j' stands for the age-period of the child. As the problem is written recursively, this implies that at every period in which parent choices affect their children's outcomes (i.e., all previous periods), the utility of their children (and future descendants) is taken into account. This formulation embeds the parental altruism motive. After the child's independence, the parent's individual problem reverts to Equation 1.

3.1.4 Retirement

At $j = 17$ (i.e., age 64), the individual retires from work (i.e., $h = 0$) and has two sources of income: savings a and publicly financed retirement benefits π . For simplicity, retirement benefits are assumed to depend on the agent's skill level. Note that individuals pay taxes on retirement benefits according to the same labor tax function $T(\cdot)$. Formally, the problem at the age of retirement is:

$$V_j(a, \theta, n) = \max_{c, a'} u(c, 0) - \bar{v}_n + \beta \widehat{V}_{j+1}(a', \theta, n), \quad (4)$$

$$c(1 + \tau_c) + \tau_n + a' - \omega - (\pi(\theta) - T(\pi(\theta))) = \begin{cases} a(1 + r(1 - \tau_a)) & \text{if } a \geq 0 \\ a(1 + r^-) & \text{if } a < 0 \end{cases}$$

$$a' \geq \underline{a}_j.$$

As in other periods, the individual can move at the beginning of each period:

$$\widehat{V}_j(a, \theta, n) = \max_{n'} \{V_j(a, \theta, n') - \bar{\kappa} 1(n' \neq n)\}.$$

3.2 Aggregate Production

We assume that there is a representative firm in each neighborhood n with the production technology $Y_n = AK_n^\alpha H_n^{1-\alpha}$, where A is the total factor productivity, K_n is aggregate physical capital in neighborhood n , and H_n is the sum of efficiency units in neighborhood n . As is standard, capital is assumed to be perfectly mobile across regions and depreciates at a fixed rate of δ per period. We assume that firms are perfectly competitive (i.e., making zero profits) and pay unit wages equal to the marginal product of labor. Formally, the equilibrium wage and return on capital are given as $w_n = (1 - \alpha)A(K_n/H_n)^\alpha$ and $r + \delta = \alpha A(H_n/K_n)^{1-\alpha}$, respectively.

As noted above, we assume individuals work in the same neighborhood n in which they reside. Since capital is freely mobile, wages are equal across neighborhoods in an equilibrium with no government intervention so our no-commuting assumption has no impact. In Section 5.2, neighborhood wages w_n will differ when we introduce the place-based wage subsidy \tilde{w}_s so that the wage in $n = 1$ is $w_1 = (1 - \alpha)A(K_1/H_1)^\alpha(1 + \tilde{w}_s)$. This type of wage subsidy programs (e.g., Opportunity Zones) tend to target those living and working in a particular area, which is in line with our assumption of people working where they live.

3.3 Housing Markets

Rental prices are determined in equilibrium given the supply functions: $S_n = \bar{S}_n \tau_n^\Delta$, where τ_n is the rent price in neighborhood n and Δ is the price elasticity of housing supply. For simplicity, we assume there are two neighborhoods denoted $n = 1$ and $n = 2$.¹⁰ Without any loss of generality, we assume that neighborhood $n = 1$ is the disadvantaged neighborhood with lower amenity (we make the normalization such that $\bar{v}_2 = 0$).

3.4 Definition of Stationary Equilibrium

The model includes J_d overlapping generations and is solved numerically to characterize the stationary equilibrium allocation. Stationarity implies an equilibrium in which the cross-sectional distribution for any given cohort of age j is invariant over time periods. Particularly important is that the distribution of initial states is determined by the choices of the older generations. The equilibrium allocation requires that households choose location, consumption, labor supply,

¹⁰Our approach is similar to prior studies. For example, Fogli and Guerrieri (2019) also construct a spatial equilibrium model that features two neighborhoods.

parental time investments, and parental transfers such that they maximize their expected utility; firms maximize profits; prices (wages, rents, and the interest rate) clear markets; and neighborhood quality s_n is equal to the total income per capita in each neighborhood.

Note that we do not require that the government budget is balanced. The government may have other non-modeled expenses G . Hence, G will be defined in the initial steady state as a residual. However, to evaluate policies (e.g., housing subsidy vouchers), we do assume that any net additional expenses must be offset by additional revenue.

3.5 Role for Government Interventions

Why might government interventions increase welfare in our model? There are two key channels. First, a main friction stems from the fact that parents cannot borrow against their child's future income. This reduces the incentive for parental investments. To illustrate this, consider a parent who is poor but pays the higher rent associated with $n = 2$ to raise a high-skilled, high-income child. This parent would want to smooth consumption intergenerationally. The fact that this rent must come at the cost of her own lifetime consumption limits their investment. If the child could promise to compensate their parent in the future (and parents were able to borrow against that future compensation), the parent would not need to reduce their consumption. Government action can make up for a parent's inability to borrow against their child's future income. Specifically, rent subsidies targeted to those with children can be thought of as (imperfectly) replacing the missing compensation mechanism via the power of taxation.¹¹ Rather than children compensating parents for their investments, the government invests directly in children and taxes them once they are adults.

Second, another key friction stems from the externalities in our model. There is inefficiency because individuals do not internalize their impact on neighborhood quality. Similar to the solution for other externality problems, the government in our context can account for the fact that individual location and work choices affect the next generation. For example, place-based wage subsidies can provide additional incentives to work in a particular neighborhood, thereby helping internalize the effect of a person's income on their neighbors.¹²

¹¹Note that there are other sources of inefficiency in the environment that motivate a government role aside from limitations on intergenerational borrowing. For example, an agent's inability to borrow fully against their own future income or to insure against future outcomes leads to imperfectly smooth consumption and worse neighborhood choices than if people were able to fully borrow. This consequence of capital market imperfections is well-understood, and a targeted rent subsidy can facilitate self-insurance and provide a lower variance of consumption. Lastly, investing in children is risky (since there are skills and wage shocks) so parents are likely to underinvest since they cannot insure against such risk. The government, instead, can pool this risk when providing rent subsidies to many families.

¹²Given that low-income families tend to live in $n = 1$, a place-based wage subsidy can also help reduce inequality (an additional outcome that a social planner may seek to change) and provide insurance against negative shocks.

Of course, any positive effects of government intervention must be weighed against the costs of increased distortionary taxation. A higher labor income tax will, *ceteris paribus*, reduce incentives to invest in human capital and work. Whether the gains outweigh the losses is a quantitative question that can be addressed using a calibrated model.

4 Estimation

This section describes how we parameterize and estimate the model. The model is estimated using simulated method of moments to match standard moments as well as more novel ones (e.g., moments informative about parental investments and the neighborhood income gap) for the U.S. in the 2000s. Some parameters can be estimated “externally,” while others must be estimated “internally” from the simulation of the model. For these, we numerically solve the steady state of this economy, obtain the ergodic distribution of the economy, and calculate the moments of interest. After estimating the model, we validate the model using reduced-form estimates from previous experimental and quasi-experimental research. The subsections below provides further details on the data and parameter estimates that we use for our calibrated model.

4.1 Preliminaries

Overview of Data and Samples: Parameters of the model are estimated to match two types of data. First, we construct individual level statistics from the following sources: the Panel Study of Income Dynamics (PSID); the National Longitudinal Survey of Youth (NLSY); and the American Time Use Survey (ATUS). Second, we use various Census data products, the ATUS, and the Opportunity Atlas from [Chetty et al. \(2018\)](#) to construct neighborhood level moments on income, housing costs, time with children, and long-run outcomes of children. The remainder of this section provides details on all data sources and the key measures we use.

Wages and the Return to Skill: The wage process and return to skills are important since they determine the career profile. We focus on a wage process that allows for differences across ages and skill levels. We propose that the wage process of a household at age j is given by the product of the wage w and efficiency units $E_j(\theta, \eta)$.¹³ These are defined as $E_j = \iota_j \psi_j(\theta, \eta)$ where ι_j is the age profile and $\psi_j(\theta, \eta)$ is the idiosyncratic component of labor productivity:

$$\log(\psi_j) = \Upsilon \log(\theta_c) + \eta_j \tag{5}$$

where θ_c is the cognitive skill level and η_j is the idiosyncratic shock. The latter is modelled as an AR(1) process: $\eta_j = \rho \eta_{j-1} + \xi_j$, where $\xi_j \sim N(0, \sigma_{\xi,j})$. An agent’s initial productivity shock η_0 is

¹³Note that due the free mobility of capital, wages are equal in the two neighborhoods so $w_1 = w_2 = w$.

drawn from a normal distribution with mean zero and variance σ_{η_0} .

We estimate this wage process in two steps. First, we estimate the age profile ι_j as a second order polynomial using PSID data. Since the model has four year periods, we estimate this income process by grouping observations over four years. We include year fixed effects (defined as the initial year of the four year period) to control for possible changes in average wages over time t and control for selection into work. We use the PSID instead of the NLSY because it includes a representative cross-section every year, so it avoids having the average age of the sample change directly with the calendar year. Specifically, we estimate the following model:

$$w_t = \beta_0 + \beta_1 Age_t + \beta_2 Age_t^2 + \beta_3 X_t + \Pi_t + \psi_t,$$

where X_t is a control for selection into work based on a Heckman-selection estimator.¹⁴ Appendix Table A1 reports the results from this estimation.

Second, we use the NLSY to identify the effect of skills on wages. We rely on the NLSY because it includes measures of skills while the PSID does not. Using the age profile estimates from the PSID data, we recover ψ_t as a residual in the NLSY data. Next, an estimate of Υ is obtained by regressing our estimate of ψ_t on the log of cognitive skills as measured by the AFQT score (i.e., we estimate Equation 5). Lastly, the AR(1) process for the residual η is estimated using the standard Minimum Distance Estimator developed by [Rothenberg \(1971\)](#). Appendix Table A1 shows the estimates obtained from our approach. These are broadly in line with those obtained in previous studies that estimate similar parameters (e.g., [Abbott et al., 2019](#); [Daruich and Fernández, 2020](#)).

Neighborhoods: As noted above, there are two neighborhoods in the model. To match this with the data, we divide U.S. Census tracts into two groups that correspond to neighborhoods $n = 1$ and $n = 2$.¹⁵ We do this in three steps. First, we use tract-level data and calculate the population-weighted percentiles of median household income within each commuting zone (CZ).¹⁶ In each CZ in the U.S., we assign all tracts that have median household income below the 10th percentile to neighborhood $n = 1$ (i.e., the disadvantaged neighborhood with low amenity value). The remaining tracts are assigned to $n = 2$. While alternative neighborhood definitions are possible, our preferred choice is motivated by the fact that real-world policies such as the MTO voucher and place-based EZ investment programs typically target the residents living in the most disad-

¹⁴To control for selection, we construct Inverse Mills ratios by estimating the participation equation using number of children as well as year-region fixed effects.

¹⁵Census tracts are small geographic units with an average population of 4,250 persons.

¹⁶Commuting zones are aggregations of counties analogous to metropolitan statistical areas. Unlike metropolitan statistical areas, commuting zones have complete coverage of the entire United States.

vantaged neighborhoods (Kling et al., 2005; Busso et al., 2013).¹⁷ Second, we compute averages of tract-level characteristics (detailed below) for the tracts assigned to $n = 1$ and $n = 2$ within each CZ. Finally, we average the statistics across CZs weighting by population.

Our approach allows us to aggregate several local area characteristics measured at the Census tract level to the two fictitious neighborhoods in our model. Table 1 reports summary statistics for the neighborhood characteristics that are the focus of our analysis. Columns 1 and 2 report summary statistics for neighborhoods $n = 1$ and $n = 2$, respectively. The percent difference between each statistic is reported in Column 3.

The following tract-level characteristics are key to our model: per capita income, home value, property taxes, and expected child outcomes. The income and housing-related measures come from the 2012-2016 American Community Survey (ACS). Note that we estimate yearly housing costs by converting home values to annual rental rates and summing this with the property tax.¹⁸ We divide this number by the average number of people in a household to obtain a per individual estimate. The measures of expected child outcomes come from Chetty et al. (2018). Specifically, we rely on tract-level statistics on the expected income for children who have parents at the 25th, 50th, and 75th percentile of the income distribution.

Table 1 shows that there are substantial differences between the less and more advantaged neighborhoods that we study. For example, the average individual income and median home values are 108 and 67 percent higher in the more advantaged neighborhood. Most importantly, the summary statistics are consistent with a model that features sorting and causal neighborhood effects: children who grow up in the more advantaged neighborhood have higher later-life incomes. This is particularly true for children from low-income households (i.e., those with parents at the 25th percentile of the national income distribution) whose incomes are 30 percent higher.

Parental Time Investment: The differences in child outcomes across areas documented in Table 1 could be due to both neighborhood effects and parental investments. Our model captures the latter by allowing for time investments. Ideally, we would estimate the relationship between parenting time and neighborhood choice directly. Unfortunately, existing time-use survey data does not provide detailed information on the neighborhood of respondents.

Due to data limitations, we rely on indirect inference to capture the pattern of parent time

¹⁷For example, the EZ program studied by Busso et al. (2013) typically aimed to target place-based investments to disadvantaged areas that contained 10 percent of the city's residents.

¹⁸The conversion is obtained by dividing home values by 20, following the literature (e.g., Fogli and Guerrieri, 2019). Housing value and property tax statistics are available at the tract and CZ-level, respectively. We impute tract-level property taxes in two steps. First, we use CZ-level data and regress property taxes on median household income. Second, we return to tract-level data and use the CZ-level regression estimates to predict property taxes based on the median household income in a given tract.

Table 1: Neighborhood Summary Statistics

	(1)	(2)	(3)	(4)
	Bottom	Top	Pct. Diff.	Area
Income				
Mean Individual Income	\$14,673	\$30,444	107.5%	Tract
Mean Household Income	\$39,348	\$81,314	106.7%	Tract
Poverty Share	28%	10%	-63.8%	Tract
Child's Mean Income at Age 26 by Parental Income				
25th Percentile	\$17,916	\$23,347	30.3%	Tract
50th Percentile	\$24,021	\$28,795	19.9%	Tract
75th Percentile	\$27,413	\$32,354	18.0%	Tract
Housing				
Median Home Value	\$150,166	\$250,378	66.7%	Tract
Property Tax	\$1,568	\$3,677	134.6%	CZ
Avg. HH Size	2.78	2.74	-1.7%	Tract
Yearly Housing Cost (Est.)	\$3,259	\$5,915	81.5%	Tract/CZ

Notes: This table reports neighborhood summary statistics for two types of neighborhoods. Columns 1 and 2 report statistics for areas that are “disadvantaged” and “advantaged”, respectively. The threshold for a disadvantaged neighborhood is based on whether median household income in the area is in the bottom 10 percent. When possible, we report summary statistics based on tract-level data. Due to data limitations, we also rely on summary statistics based on commuting zone (CZ) level data. CZs are geographical aggregations of counties that are similar to a metro area but cover the entire United States (including rural areas). We convert the CZ-level statistics to tract-level measures to match the two neighborhoods in our model. Yearly housing costs are estimated combining tract-level home values and CZ-level statistics on property taxes, as explained in the main text. Housing, demographic, and income statistics are from the 2012-2016 ACS. Child outcome statistics are from the Opportunity Atlas [Chetty et al. \(2018\)](#).

investments across neighborhoods. Our approach is based on two steps. First, we aim to estimate the following model that relates time investment to household income:

$$\log(\tau) = \varrho_0 + \varrho_1 \log(y) + \varrho_2 X + \epsilon,$$

where X is a set of controls variables for the respondent’s household characteristics, the age of the youngest child, and the total number of children in the household. Our main interest is the estimate of ϱ_1 which is the elasticity of parental time with respect to family income.

Second, we set the estimate of the elasticity ϱ_1 as a key moment for our calibration of the parameters of the child skill production function (detailed further below). In particular, we use this moment to estimate the complementarity between parenting time and neighborhood quality (i.e., γ). Intuitively, the elasticity that we estimate will be larger if there is stronger complementarity

between parenting time and neighborhood quality because higher income parents tend to live in more advantaged neighborhoods. Our aim is to match the relationship between parenting time and income to address concern over the risk of misattributing the impact of parent time investments as neighborhood effects.

To estimate the relationship between parental time investment and family income, we rely on the ATUS.¹⁹ We create a sample of parents and measure the amount of “quality time” spent with each child in the household. The sample of parents includes all ATUS respondents surveyed during the period 2003-2019 and who were ages 18-65 and had at least one child in the household.²⁰ We follow [Price \(2008\)](#) and define quality time to include all activities in which either the child was the primary focus of the activity or in which there would be a reasonable amount of interaction (e.g., eating together). We compute the total time that a parent spends with their children and construct a per capita (child) measure by dividing by the number of children in the household. We scale this measure by two when the respondent has a partner in the household to obtain a measure of the average amount of parental time that a child receives. Based on this ATUS sample, Appendix Table [A2](#) shows that we consistently estimate that ϱ_1 is equal to 0.10 across a range of specifications.

Child Skill Development: We estimate children’s future skills as being dependent on current skills, parental skills, and an index of investments. Investments are a function of neighborhood income and parental time inputs. As explained in Section 3, we assume that the child development function has a nested CES form. For the outer CES, we use estimates of the parameters from [Cunha et al. \(2010\)](#), which are based on a representative sample (see Appendix [B1](#)). These estimates are specific to age-period j (i.e., the parameters vary with the child’s age). A key finding from their work is that skills are more malleable when children are young (i.e., the elasticity of substitution determined by $\rho_{q,j}$ is larger for younger children). We also follow [Cunha et al. \(2010\)](#) in assuming that skills are a vector of both cognitive and non-cognitive skills. In their work, [Cunha et al.](#) highlight that failure to allow for these two types of skills leads to estimates that

¹⁹The ATUS sample is based on a group of households in the outgoing rotation of the Current Population Survey (CPS). For each household, one adult is randomly selected to complete a detailed survey on the respondent’s activities of the previous day. For each activity, they report the starting and ending time, their location, and the members of their household who were present during the activity. Because the ATUS respondents also participated in the CPS, we have detailed information on household characteristics such as family income.

²⁰Descriptive statistics for our ATUS sample are as follows: the average age was 39 years old, and the average number of children is 1.91. Note that we restrict the sample to respondents with non-missing geographic information. County-level residential information is suppressed when a respondent resided in an area where the population was less than 100,000. We use information on county of residence to assign individuals to a CZ. As a robustness check, our most stringent specification estimating the elasticity of time with respect to income controls for CZ fixed effects. Note that our elasticity estimates are similar when we do not restrict the sample based on having county information.

suggest investments for low-skill children are much less productive. Thus, θ and θ_k are vectors with an entry for each skill type.

The initial draw of skills is assumed to depend on parent skill as an AR(1) process that is independent for cognitive and non-cognitive skills:

$$\log(\theta_{k,q}) = \hat{\rho}_q \log(\theta_q) + \epsilon_{\theta_{k,q}}, \quad q \in \{c, nc\},$$

where $\epsilon_{\theta_{k,q}}$ is a shock, independent across skill types. We define the persistence component $\hat{\rho}_q$ to be equal to $\rho_k \times \left[\frac{\text{Var}(\log(\theta_{k,q}))}{\text{Var}(\log(\theta_q))} \right]^{0.5}$, where ρ_k is the correlation between $\log(\theta_{k,q})$ and $\log(\theta_q)$. We use estimates of the variance terms directly from [Cunha et al. \(2010\)](#) to calculate $\hat{\rho}_q$. Note that the variance of the skill shock is given by: $\epsilon_{\theta_{k,q}} = \text{Var}(\log(\theta_{k,q})) - \hat{\rho}_q^2 \text{Var}(\log(\theta_q))$.

We assume the following functional form for neighborhood spillovers. As mentioned above, the neighborhood quality effect on children is summarized by the sum of capital and labor income per capita: $s_n = \bar{y}_n + (r + \delta)\bar{a}_n$, where the terms \bar{y}_n and \bar{a}_n are labor income and asset holdings per capita in neighborhood n . For the functional form of neighborhood effects, we assume that $f(s_n) = A s_n^\zeta$. Intuitively, a larger value for the parameter ζ allows neighborhoods to have a larger impact on child development.

In this framework, there are three sets of parameters governing investments. We internally estimate the parameters $\alpha_{I,j}$ and ζ to match two key moments for the average difference in child outcomes between advantaged and disadvantaged neighborhoods: the difference in average incomes for low-income children (i.e., have parents at the 25th percentile of the income distribution) and the differences in average incomes for high-income children (i.e., have parents at the 75th percentile of the income distribution).²¹ We estimate the neighborhood parameter γ to match the elasticity of parental time to parental income (as discussed above), which implies that the associated elasticity (i.e., $1/(1-\gamma)$) is 0.49. Although not directly comparable, this elasticity is within the range of existing estimates for the complementarity between parental time investments and other educational inputs in children's development (e.g. [Caucutt et al., 2020](#); [Abbott, 2021](#)).²²

Taxes, Lump-sum Transfers, and Pension Benefits: Our model features several margins of taxation. For the labor income tax function, we assume that $T(y) = y - \lambda y^{1-\tau_y}$. The parameter τ_y helps determine the progressivity of the marginal tax rate. We use the preferred estimate of $\tau_y =$

²¹Note that we allow $\alpha_{I,j}$ to vary between the periods when children are the youngest ($j = 8$) and subsequent age-periods ($j = 9 - 11$). This captures the fact that parental time investments are decreasing with the age of children. Since $\alpha_{I,j}$ varies, we similarly allow the investment scaling parameter \bar{A}_j to take different values between the periods when children are the youngest ($j = 8$) and in the later periods ($j = 9 - 11$).

²²Note that, without loss of generality, we set the scaling parameter A such that the quality of neighborhood $n = 2$ is normalized to one (i.e., $f(s_2) = 1$) in the baseline steady state.

0.18 from [Heathcote et al. \(2017\)](#). We estimate λ to match a labor income tax revenue equal to 22.4 percent of gross labor income [McDaniel \(2007\)](#).²³ In addition to labor taxes, the government taxes consumption and capital income. Based on [McDaniel \(2007\)](#), we set $\tau_a = 0.266$ and $\tau_c = 0.079$.

The model also features a lump-sum transfer ω that we estimate to match a measure of income redistribution—the ratio of the variance of pre-tax total (i.e., labor and savings) income to after-tax total income—to capture the disposable income available at the bottom of the income distribution. We find that $\omega = \$2,425$ on an annual basis. Note that lump-sum transfers are a standard feature in equilibrium models such as ours. The justification for this stems from the observation that low-income households tend to have higher after-tax income than what would be predicted based on a tax function without a lump-sum component (e.g., Figure 1 of [Heathcote et al., 2017](#)).

Finally, our model features pension benefits based on the Old Age, Survivors, and Disability Insurance U.S. federal program. As detailed in Appendix B2, we use skill levels in the model to estimate the average lifetime income on which the replacement benefit is based.

Preferences: We specify that the period utility function for consumption and labor is:

$$u(c, h) = \frac{c^{1-\sigma_c}}{1-\sigma_c} - \mu \left(\frac{h^{1+\theta_h}}{1+\theta_h} \right),$$

where we follow the literature and specify $\sigma = 2$ and $\theta_h = 3$ (implying the Frisch elasticity is one-third). We estimate the scaling parameter μ to match the average number of working hours observed in the PSID data. When parents choose their time to spend with children, we assume the disutility is assumed to be linear: $v(\tau) = \xi\tau$. The parameter ξ is estimated to match the average time spent with children. Finally, the altruism factor $\tilde{\beta}$ in Equation 3 is estimated to match the average monetary transfer from parents to children in the PSID.²⁴

Prices: Wages are normalized such that the average annual income at age 48 is equal to 1 in the model. In the PSID data, this income is equal to \$36,575.

Aggregate Production Function: As noted above, we assume there are there is representative firm in neighborhood n with the production technology $Y_n = AK_n^\alpha H_n^{(1-\alpha)}$. We set $\alpha = 1/3$. Capital depreciates at rate $\delta = 0.065$.

Housing Markets: Rental prices are determined in equilibrium given the supply functions: $S_n = \bar{S}_n \tau_n^\Delta$, where τ_n is the rent price in neighborhood n and Δ is the price elasticity of housing supply.

²³The estimated economy has an average marginal income tax rate of 38 percent.

²⁴We follow the steps in [Daruich \(2020\)](#) in our sample. We estimate the average total transfers received by children when they are aged 17–26 to be \$40,837 per child, equivalent to 125 percent of average annual individual income. The transfer data include small and large (e.g., to buy houses or cars) transfers, in-kind transfers (i.e., college tuition), and estimates for housing costs if the child lives with the parents. See cited paper for details.

The standard estimate in the literature comes from Saiz (2010). They find that the population-weighted average price elasticity in the U.S. is 1.75, so we set $\Delta = 1.75$. Based on that elasticity and our neighborhoods definition such that 10 percent of people live in $n = 1$, we can back out \bar{S}_1 and \bar{S}_2 using the housing costs reported in Table 1. This leads to $\bar{S}_1 = 11.9$ and $\bar{S}_2 = 37.6$.

4.2 Simulated Method of Moments: Results

As previewed above, there are 15 parameters of the model that we estimate. We use simulated method of moments to estimate the following parameters: μ , $\tilde{\beta}$, \bar{v}_1 , $\bar{\kappa}$, σ_κ , ξ , γ , $\bar{A}_{j=1}$, $\bar{A}_{j\neq 1}$, $\alpha_{l,j=1}$, $\alpha_{l,j\neq 1}$, ρ , \hat{A} , λ , and ω . Specifically, we use a Sobol sequence in order to solve and simulate the model in a fifteen-dimensional hypercube in which parameters are distributed uniformly and over a “large” support. This provides a global method to find combinations of parameters.

Table 2: Estimation Parameters and Moments

Parameter	Value	Description	Moment	Data	Model
Preferences					
μ	307.9	Mean labor disutility	Avg. hours worked	32.9	32.2
$\tilde{\beta}$	0.32	Altruism	Parent-to-child transfer as share of income	125.4%	128.8%
Neighborhood Value and Moving Costs					
\bar{v}_1	0.41	Exogenous disutility of $n = 1$	Income neighborhood ratio	107.5%	104.6%
$\bar{\kappa}$	3.47	Moving cost	Share in $n = 2$	90.0%	89.3%
σ_κ	1.29	Moving cost shock	Share of young ($j = 5 - 7$) in $n = 2$	85.6%	85.6%
Skill Formation: $I = \bar{A} [\alpha_{l,j} f (\bar{y}_n)^\gamma + (1 - \alpha_{l,j}) t^\gamma]^{1/\gamma}$					
ξ	0.17	Parent disutility of time with children	Avg. weekly hours with child (age 0-3)	25.6	24.8
γ	-1.05	Neighborhood-time substitutability	Reg. of Log-Time on Log-Income	0.10	0.10
$\bar{A}_{j=1}$	3.41	Returns to investments ($j = 1$)	Average log-skills (age 4)	0.00	0.00
$\bar{A}_{j\neq 1}$	2.90	Returns to investments ($j = 2 - 4$)	Average log-skills (age 16)	0.00	-0.01
$\alpha_{l,j=1}$	0.34	Neighborhood share ($j=1$)	Child inc. diff.: 25th pct. parents	30.3%	30.3%
ρ	3.64	Neighborhood curvature	Child inc. diff.: 75th pct. parents	18.0%	18.2%
$\alpha_{l,j\neq 1}$	0.94	Neighborhood share ($j = 2 - 4$)	Avg. weekly hours with children (age 4-15)	15.2	17.5
\hat{A}	1.25	Neighborhood scaling	Neighborhood $n = 2$ normalization	1.00	1.00
Government					
λ	0.74	Tax function scalar	Aggregate labor income tax rate	22.4%	22.5%
ω	0.05	Lump-sum transfer	Income variance ratio: Disposable to pre-gov	0.61	0.62

Notes: This table reports estimates of the model parameters as well as the observed and simulated moments associated with each parameter estimate. See text for definitions and data sources.

Table 2 reports estimated parameters as well as the corresponding moments in data (Column 5) and the simulated economy (Column 6). Overall, the model provides a good fit of the data. Given our purposes, we highlight that the simulated moments related to the skill formation parameters

are close to their empirical counterparts. Moreover, the simulated moments that are informative for the neighborhood value parameter and costs are also close to the ones observed in the data.

4.3 Validation Exercises

4.3.1 Comparing Simulations with Experimental Estimates of Housing Voucher Effects

We begin our validation exercises by using credible estimates from the literature to test the most important and novel feature of our model: the influence of neighborhoods on child development. [Chetty et al. \(2016\)](#) studied the Moving to Opportunity (MTO) experiment which provided housing vouchers to low-income families living in impoverished neighborhoods in Baltimore, Boston, Chicago, Los Angeles and New York. Families were randomized into one of three groups: an “experimental” group, a Section 8 comparison group, and a control group. The experimental group received housing vouchers that could only be used to subsidize rent for private market housing located in Census tracts with poverty rates below 10 percent. Families in the Section 8 Comparison group received vouchers that could be used without any neighborhood restrictions. Members of the control group received no vouchers through this experiment.

Prior studies of MTO have found that the program reduced the likelihood of living in a high-poverty neighborhood and had beneficial impacts on long-run outcomes of children. [Chetty et al. \(2016\)](#) find that moving through MTO increased earnings of children who moved by \$3,500. The pattern in the MTO results is consistent with a model of childhood exposure effects in which exposure to “better” environments leads to improved long-run outcomes.

We simulate a policy similar to the MTO voucher program using our model. From the steady state, we evaluate a scenario where the government provides low-income families that have children and live in the disadvantaged neighborhood with a voucher that subsidizes rent for housing in the more advantaged area. In our simulation, we limit eligibility to individuals with incomes below the tenth percentile of income.²⁵ The subsidy in our simulation covers 100 percent of rent and must be redeemed in the advantaged neighborhood.²⁶ Note that this validation exercise also assumes that rental prices and other equilibrium quantities (such as neighborhood quality) do not change. These assumptions are in line with the idea that relatively few families move in a small-scale RCT such as MTO, implying that neighborhood characteristics are not affected.

Voucher-eligible families make two key choices in our model. First, they must decide whether to take-up the voucher and relocate to the more advantaged neighborhood. We find that 70 per-

²⁵Results from our validation exercise are similar when we use alternative low-income definitions.

²⁶The standard U.S. voucher program requires households to pay 30 percent of their income as a rental contribution. While our model does not include this feature, we suspect it may have little bearing on our validation exercise since 73 percent of MTO households were unemployed and thus did not pay a rental contribution ([Katz et al., 2001](#)).

cent of households opt for the voucher in our MTO-based simulation.²⁷ This is broadly comparable with the take-up rates for MTO households with young children in the experimental and Section 8 voucher groups (48 and 66 percent, respectively). Second, parents adjust the amount of time that they spend with their children. Given that time and neighborhood quality are complements in our estimation, parents with young children ($j = 8$) spend on average 6.4 (60 percent) more weekly hours with their children after taking-up a housing voucher.

Our main finding is that the voucher-subsidy program in this simulation exercise generates similar positive impacts on long-run outcomes of children. We calculate that children in our simulation have 30.8 percent higher income when children are in their late 20s. This effect can be compared to the MTO results in two ways. First, this simulated impact is nearly identical to the average 31 percent increase in earnings experienced by children whose household moving using an MTO experimental voucher. Second, we can also compare the simulated effect on earnings to the range of effects the site-specific treatment effects observed in the MTO demonstration.²⁸ Figure 3 plots dots (in black) for the treatment effects for the unrestricted (i.e., standard Section 8) and experimental voucher groups in each of the five MTO cities. The results show that reductions in neighborhood poverty were larger for treated households in the experimental group. In line with this, the treatment effects on the earnings of children generally increase with the larger improvements in neighborhood quality (i.e., reductions in neighborhood poverty rates). The dashed line plots the predictions from a linear regression through the site-specific estimates. The diamond (blue) point on the figure represents the results from our simulation.²⁹ Reassuringly, we see that the simulation generates results that are very close to the simple linear prediction based on the pattern of site-specific MTO results.

4.3.2 Comparing Simulations with Quasi-experimental Estimates of Place-Based Policy Effects


As a supplemental validation exercise, we test whether simulations from our calibrated model match credible reduced form evidence on the effects of place-based policies on labor markets. [Busso et al. \(2013\)](#) studied the impact of the EZ program, one of the largest federal policies in the U.S. that provided incentives to encourage development of distressed and economically underperforming areas. A key feature of the policy was a wage subsidy. Specifically, firms operating

²⁷This margin is affected by the eligibility restrictions for the program. A more restrictive voucher program (e.g., limiting eligibility to those in the bottom 5 percent of income rather than bottom 10 percent) would lead to lower take-up because moving costs make potential beneficiaries concerned about the prospect of losing the voucher benefit in the future if their income is higher.

²⁸Site-specific treatment effects on long-run earnings of children and household poverty rates are from [Chetty et al. \(2016\)](#) and [Ludwig et al. \(2013\)](#).

²⁹Neighborhood poverty in the model is based on the two fictional neighborhoods and the associated averages of the tract-level Census data (as described in Section 4.1).

Figure 3: Validation: Comparing MTO Site-Specific Effects to Model Simulation



Notes: This figure compares results from the Moving to Opportunity (MTO) housing voucher experiment and simulation results based on our calibrated model. The MTO experiment took place in five cities: Baltimore, Boston, Chicago, Los Angeles and New York. Families were randomized into one of three groups: an experimental group, a Section 8 comparison group, and a control group. The solid (black) dots plot the treatment effects on long-run earnings of children (y -axis) and the change in neighborhood poverty for each site and voucher group (x -axis). Since there were five sites and two vouchers arms (Section 8 and Experimental, respectively), there are 10 solid dots. The solid (blue) diamond plots the simulated effect on long-run earnings of children and associated change in poverty rates from moving to the more advantaged neighborhood ($n = 2$). The dashed line shows predictions from a linear regression of the treatment effects on long-run child income (as a percent effect relative to the control group mean) on the reductions in poverty rates.

in a designated EZ became eligible for an employment tax credit of up to 20 percent of the first \$15,000 in annual wages earned by an employee who lived and worked in the area.

Busso et al. (2013) found that the EZ program had large positive impacts on labor market activity. To estimate effects, they used a difference-in-differences approach that compared Census tracts selected for the EZ program to a comparison group of tracts that had applied but been rejected for EZ designation. Their results (see Panel C of Table 6 of their study) indicate that the program increased annual wage income in 2000—approximately five years after the program began—by 17 to 24 percent for local residents who worked in the designated area.

We simulate a policy experiment to mirror the EZ program using our calibrated model. Relative to the steady state with no intervention, we study a scenario in which the government provides a wage subsidy of 20 percent to all workers who live (and by assumption work) in the disadvantaged neighborhood $n = 1$ for a specified duration.³⁰ In this simulation, we allow rental prices and neighborhood quality to change. These assumptions are in line with the idea that the EZ designation in Busso et al. (2013) was sufficiently large to have equilibrium impacts.

³⁰Note that the EZ program did not provide permanent employment tax credits. Tax credits were available to a business for as long as ten years.

Our main finding is that the simulated impacts of a 20 percent subsidy on income are similar to the positive impacts found in [Busso et al. \(2013\)](#). Table 3 shows a range of simulated impacts where each column varies the duration of the 20 percent wage subsidy from 1 to 4 periods, and the rows report impacts on income in different periods. For example, we find that adults in $n = 1$ have 21 percent higher earnings when the program is only run for one-period (corresponding to four years). Although the effects vary slightly with the intervention duration and time of evaluation, the range of estimates is broadly comparable to [Busso et al. \(2013\)](#).

Table 3: Validation: Simulated Impacts of a 20-percent Wage Subsidy on Income

	(1)	(2)	(3)	(4)
	$\Delta \text{Log-income}$			
Period of Evaluation				
First	0.21	0.19	0.18	0.18
Second	-0.01	0.26	0.25	0.25
# of Subsidy Periods	1	2	3	4

Notes: This table presents results from simulations of a 20-percent wage subsidy provided to residents of the disadvantaged neighborhood $n = 1$ using our calibrated model. Each column reports simulated effects (relative to a steady state with no government intervention) on the income of $n = 1$ residents when the wage subsidy program lasts for one, two, three, or four periods. The rows report results on income evaluated in the first or second periods.

4.3.3 Comparing Child Outcomes Using Alternative Neighborhood Definitions


Our calibrated model matches the differences in average incomes for children who grew up in neighborhoods with different levels of average income (controlling for parental income). Specifically, the model is based on comparisons between children who grew in neighborhoods with average income in the bottom 10 percent to those who grew up in the top 90 percent. While this simple grouping of neighborhoods allows us to map the usual targets of Moving to Opportunity and Empowerment Zone programs, we can also exploit variation within the top 90 percent of neighborhoods to further assess if our model’s predictions of the impact of neighborhood quality on child outcomes are in line with data.

As an additional validation test, we study a partial-equilibrium exercise in which we modify the definition for neighborhood $n = 2$ in our model using various neighborhood income deciles. In particular, we modify the neighborhood quality and rental costs in our model to match the observed average income per capita and living cost for neighborhoods belonging to each decile between two and ten.³¹ With the exception of the neighborhood amenity parameter (which changes

³¹In our baseline estimation, the advantaged neighborhood’s average income is between that of deciles 6 and 7.

to ensure that the share of people in $n = 1$ remains unchanged), all other parameters remain at their originally estimated values.

Figure 4: Validation: Child Outcome Differences Under Alternative Neighborhood Definitions



Notes: This figure shows the results of partial-equilibrium exercises that alter the advantaged neighborhood ($n = 2$) definition in our model. This exercise creates nine versions of the model in which we set neighborhood quality and rental cost based on the neighborhoods belonging to each decile between two and ten of distribution of neighborhoods. The neighborhood amenity parameter is adjusted to keep the share of people in $n = 1$ unchanged across simulations. All other parameters are kept at their originally estimated values reported in Table 2. The figure reports the model's predictions on children income differences (in squares) with those observed in the data (in diamonds) separately for those with low-income (blue) and high-income (purple) parents.

Figure 4 compares predicted child income differences from each of the alternative versions of our model to the data. Each square is the difference between the top and bottom neighborhood from a simulation, and the diamonds are the corresponding empirical moments based on the Opportunity Atlas. We provide separate results for the children with parents at the 25th (blue) and 75th (purple) percentile of the income distribution. Notably, the model predictions appear to match the large differences in outcomes that we observe as the neighborhood quality difference increases. In line with the data, the model also predicts that the differences are larger for children of low-income parents, particularly so when the neighborhood quality differences are larger. We interpret these finding as suggesting that our calibrated model is able to capture the effect of various levels of neighborhood quality differences on child outcomes. This provides confidence that our model predictions are not limited to studying the type of neighborhood quality changes captured in our MTO-based validation exercise in Section 4.3.1.

5 Policy Analysis

This section quantitatively evaluates the general equilibrium effects of government interventions that change neighborhood quality for children. As in most OLG models with intergenerational human capital investments, the main rationale for government involvement, as detailed in Section 3.5, stems from the fact that children do not control the inputs into their development or compensate their parents for doing so. This can lead to reduced levels of childhood investment relative to an economy where it is possible for parents and children to sign contracts to facilitate intergenerational transfers.³² In addition, the presence of externalities in our environment provides another motivation for intervention since government policies can help internalize the effect of a person’s income on other individuals.

As previewed above, our analysis focuses on two types of interventions: housing rental vouchers and place-based wage subsidies. For each program, we begin our analysis by evaluating the large-scale and long-run effects of several alternative versions of the program that vary programmatic features such as the subsidy rates. We start by identifying the version of each program that has the highest steady-state welfare gains. Note that welfare is defined by consumption equivalence under the veil of ignorance.³³ We compare the highest steady-state welfare policy to several alternatives in order to understand the quantitative importance of each feature of the voucher or wage-subsidy programs. Focusing on the welfare-maximizing policy, we provide a decomposition analysis to study the role of the different equilibrium forces (e.g., long-run intergenerational dynamics, limited housing supply, endogenous neighborhood quality, and taxation) and study transition dynamics.³⁴ We conclude this section with a discussion of the distribution of each program’s impact on welfare and how this may have political economy implications.

5.1 Evaluating Alternative Housing Voucher Programs

This section studies the consequences of a housing voucher program after accounting for the general equilibrium forces that are difficult to evaluate using existing empirical evidence based on RCTs and small-scale natural experiments. The policy is such that individuals in the program would pay $\tau_2 \times (1 - s)$ if living in $n = 2$, where s is the rent subsidy rate. The government finances these programs by adjusting the labor tax parameter λ (which is typically thought of as governing the average labor tax rate) such that its budget is unchanged from the initial steady state.

³²Previous theoretically focused research on parent investment highlights that childhood development policies (e.g., Loury, 1981; Baland and Robinson, 2000) and investments in neighborhood or schooling programs (e.g., Benabou, 1996a; Durlauf, 1996; Fernandez and Rogerson, 1996) can be welfare enhancing.

³³See Appendix C for details on the definition of our welfare measure.

³⁴In our transition analysis, we study the impacts on welfare for adults alive at the time of a policy’s introduction.

To simplify the analysis, we focus on the program that generates the highest steady state welfare (under the veil of ignorance). To determine this, we simulate versions of our model which differ in terms of three characteristics: (1) the voucher subsidy rate s ; (2) an eligibility restriction in terms of the individuals hourly wage $wE_j(\theta, \eta)$; and (3) an eligibility restriction based on the presence of children (which is based on age j given the exogenous fertility assumption in our model). We search over 50 variations of voucher programs defined by these characteristics.³⁵ After determining the policy with characteristics that leads to the largest steady-state welfare gains (relative to the baseline scenario where there is no voucher program), we similarly report results for several alternative policies to better understand which policy features drive the gains.

The highest steady-state welfare gains are achieved with a policy that has a full subsidy rate and targets households that have children and wages below the ninth decile (i.e., the 90th percentile). In this scenario, the share of children living in the more advantaged neighborhood increases by 12.9 percent.³⁶ As lower-income parents and more children move to $n = 2$, average income per capita is reduced in this location, leading to a reduction in its neighborhood quality of 3.8 percent. Despite this, the average neighborhood quality to which children are exposed to increases by 2.2 percent. This leads to a 1.4 percent increase in labor productivity (calculated by $\psi = e^{\gamma \log(\theta_c)}$ as detailed in Equation (5)). As more people are willing to move to $n = 2$, the rent τ_2 increases by 3.5 percent. Even though rent increases and funding the policy requires a large increase in taxes, with the marginal tax rate increasing by 16.0 percent, the positive effects of the policy dominate and welfare increases by 4.3 percent.

Table 4 reports results that permit a comparison of the policy that achieves the highest steady-state welfare gains to alternatives. Columns 2 and 3 show that voucher programs which reduce the subsidy rate to 10 and 90 percent decrease take-up of the voucher program and the share of children in $n = 2$ increases by only 0.5 and 11.8 percent, respectively. These increases are smaller than the 12.9 percent increase observed in the highest welfare policy. Thus, the welfare gains are relatively smaller at 0.2 and 4.0 percent in these scenarios. Columns 4 and 5 show that

³⁵We allow for five different subsidy rates (i.e., 20, 40, 60, 80, and 100 percent). We target the levels of wage targeting by quintiles of the cross-section distribution (e.g., below the 1st quintile or below the 3rd quintile). Finally, targeting by the presence of children means that only those age $j = 8 - 11$ can obtain the subsidy.

³⁶The subsidy take-up rate in the long-run optimal scenario is nearly 100 percent. While this number is larger than the typically finding in empirical studies of rent vouchers (Chyn et al., 2019), it is difficult to compare our long-run simulation with empirical estimates. Individuals in the long-run analysis are likely to be born in and stay in the more advantaged neighborhood. In contrast, empirical evidence on vouchers take-up rates typically studies households *already* living in disadvantaged areas who may face significant moving costs. With this in mind, we believe the validation exercise from Section 4.3 is a reasonable comparison. Our validation aims to replicate findings from the MTO experimental voucher study by studying a short-run program targeted to families living in the disadvantaged neighborhood. This exercise finds a lower take-up rate of 70 percent—much closer to empirical estimates.

Table 4: Long-Run Effects of Alternatives Housing Voucher Programs

	(1)	(2)	(3)	(4)	(5)	(6)
	Long-Run Optimal	Alt. Subsidy Rates		Alt. Wage Targets		Alt. Demo.
Rent Subsidy						
Subsidy Rate	100%	10%	90%	100%	100%	100%
Wage Below Decile	9	9	9	2	10	9
Target Children	✓	✓	✓	✓	✓	✗
Welfare	4.35%	0.21%	4.04%	2.46%	4.34%	-8.00%
Policy						
ΔTax Rate	16.0%	1.5%	14.3%	2.9%	16.5%	44.1%
Share with subsidy	24.3%	24.3%	24.3%	4.9%	25.0%	87.6%
ΔShare of children in $n = 2$	12.9%	0.5%	11.8%	2.4%	12.9%	13.0%
Aggregates						
ΔGDP	-0.6%	-0.2%	-0.5%	-0.3%	-0.5%	-6.7%
ΔCapital	-1.3%	-0.3%	-1.1%	-1.6%	-1.1%	-13.9%
ΔLabor Productivity	1.4%	-0.0%	1.3%	0.4%	1.5%	-0.9%
Neighborhoods						
ΔQuality $n = 1$	-1.7%	-0.6%	-4.4%	12.7%	-1.9%	39.5%
ΔQuality $n = 2$	-3.8%	-0.3%	-3.4%	-2.0%	-3.8%	-11.8%
ΔQuality for avg. child	2.2%	-0.1%	2.1%	-0.2%	2.2%	-6.2%
ΔWage $n = 1$	-0.4%	-0.0%	-0.3%	-0.7%	-0.3%	-3.9%
ΔWage $n = 2$	-0.4%	-0.0%	-0.3%	-0.7%	-0.3%	-3.9%
ΔRent $n = 2$	3.5%	0.3%	3.4%	1.3%	3.5%	6.1%

Notes: This table presents results from an analysis of the effects of various housing voucher programs on equilibrium outcomes. All effects are calculated by comparing the difference in outcomes between a simulation for a given voucher and the baseline scenario where there is no government housing intervention. Column 1 reports differences when comparing the policy that generates the highest welfare gain relative to the baseline scenario. Columns 2-3 presents results for scenarios where the voucher subsidy covers 10 and 90 percent of rent, respectively. Column 4 presents results for scenarios where the voucher eligibility targets those whose hourly wages, $wE_j(\theta, \eta)$, are below the second decile. Column 5 abstracts from wage targeting. Column 6 presents results for the scenario where the voucher eligibility does not depend on whether the household has children.

restricting eligibility based on parent earnings has limited impacts on take-up and corresponding small differences in welfare gains. For example, targeting the program to only those with wages below the second decile still leads to welfare gains of 2.5 percent. Notably, it is cheaper to finance this policy so taxes only need to increase by one-fifth as much (i.e., 2.9 versus 16.0 percent).

This may matter for the political economy of a voucher program, a point we return to when we discuss transition dynamics. A universal policy for families with children (Column 5) leads to very similar effects as the long-run optimal policy. This result implies that wage targeting has little impact on our policy conclusion.³⁷ Finally, the last column reports steady-state outcomes in a scenario which mirrors the policy that delivers the highest welfare but eliminates the requirement that individuals have children. Under this less targeted program, almost everyone lives in $n = 2$. This leads to increases in both rent prices and average marginal tax rates of 6.1 percent and 44 percent, respectively. These tax increases lead to sizable reductions in labor supply, thereby lowering income and notably decreasing neighborhood quality in $n = 2$. These negative effects sum up to large welfare losses of 8.0 percent, suggesting that targeting families with children is crucial for the rent-voucher program to have positive effects.

To summarize, the comparison across columns of Table 4 suggest there are two crucial components of the rent voucher program. The program should feature a high subsidy rate such that the take-up is high and restrict eligibility only to the families with children to avoid the large increases in taxes that have large negative effects on income and neighborhood quality. While wage targeting does have an effect on welfare, its importance appears to be much smaller than the other two program parameters.³⁸

5.1.1 *Voucher Decomposition*

The results so far demonstrate that a large-scale targeted housing voucher program can notably increase long-run welfare. This section provides an analysis that traces out the mechanisms that drive these gains. We define the highest welfare achieving policy introduced in the previous section as our benchmark and compare this to simulations that shut down several equilibrium channels in our model.

Table 5 presents decomposition results for our main economic outcomes. The bottom row in bold reports statistics for the benchmark, including all equilibrium effects. The first row reports simulation results when the voucher program defined in our benchmark simulation is introduced for one generation (starting from the initial steady state without a rent-voucher program), there are no general equilibrium effects (i.e., rents, neighborhood quality, wages, and interest rates do

³⁷Our approach to targeting focuses on wages because this simplifies the computational tractability of our analysis. A more realistic approach would consider targeting by income. A benefit of modeling income-based targeting is that our analysis would better capture distortions in labor supply induced by vouchers. That said, the similarity between our optimal and universal voucher results suggests that considering income-based targeting will not ultimately affect the voucher policy that we study for our welfare analysis.

³⁸Less wage-targeting leads to higher take up, increasing gains, but also leads to higher tax increases, reducing the gains. These two opposing effects make the welfare effects of rent voucher programs relatively constant across different levels of wage targeting.

not change), and the economy has no requirement to balance the government’s budget. Effects are evaluated for the children of the single generation that is offered the voucher. This initial simulation is intended to be representative of what we would expect from an RCT, which is typically implemented on a small scale (i.e., taxes, prices and neighborhood qualities do not change) and applied to only to members of one generation.³⁹

Table 5: Rent Subsidy: Decomposition of Equilibrium Forces for Main Economic Outcomes

Equilibrium Forces					Change from Initial Steady State (%)						
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Long Run	Taxes	Real Estate	Neigh. Qual.	Prices	Share in $n = 2$	Neigh. Qual, $n = 2$	Income Ineq.	Upward Mobility	Labor Prod.	Welf. Adults	Welf. Children
X	X	X	X	X	0.8	0.0	-7.3*	27.9*	0.9*	1.8	5.7*
X	✓	X	X	X	0.7	0.0	-7.0*	28.0*	0.8*	0.4	4.4*
✓	✓	X	X	X	11.1	0.0	-6.6	28.0	2.3	-	6.6
✓	✓	✓	X	X	6.8	0.0	-2.2	28.1	2.3	-	5.3
✓	✓	✓	✓	X	6.8	-4.5	-2.6	30.0	1.3	-	3.3
✓	✓	✓	✓	✓	7.0	-3.8	-2.7	29.4	1.4	-	4.3

Notes: This table presents results that decompose how various equilibrium forces affect economic outcomes under the housing voucher program that generates the highest steady-state welfare gains. Each row provides results from a separate simulation of the model. Columns 1-5 describe which equilibrium force is shut down in each simulation. Columns 6-12 report changes (in percent) in outcomes calculated by comparing each simulation to the initial steady state (where there is no housing voucher program). The asterisks highlight that short-run effects are evaluated for children of the cohort that received the policy intervention.

The main finding apparent in this section is that the welfare gains predicted based on an RCT evaluation of the benchmark program are broadly similar to those obtained in the long-run steady state. Column 12 shows that welfare (as measured by the children) increases by 5.7 percent under a small-scale voucher program in the short-run. The 4.3 percent increase that is produced in the long-run benchmark, as introduced in the previous section, is about 25 percent smaller. In other words, the welfare improvements predicted from a small-scale voucher program are only partially offset by the equilibrium effects captured in our model.

The following rows provide results that help unpack the role of the equilibrium forces. The second row takes into account that taxes need to increase to finance the policy. Note that we hold constant several other endogenous variables such as housing market prices, neighborhood quality, labor market conditions, and other equilibrium objects (e.g., wages and interest rates).

³⁹Note that the initial simulation in the top row of Table 5 differs from the validation simulation in two main ways. First, the simulation for the validation exercise only targets those with wages below the tenth percentile. In contrast, the simulation in this section sets voucher eligibility to those with wages below the fourth quintile. Second, the validation simulation only considers providing vouchers to individuals who live in the disadvantaged neighborhood initially. In this section, the simulation in the top row provides vouchers to anyone living in the good neighborhood who meets the eligibility conditions.

As measured by the children, the welfare gains are reduced by 1.3 percentage points (about 25 percent of the gains without equilibrium forces of 5.7). The third row allows for long-run effects: the benchmark voucher program is implemented permanently and effects are evaluated in the new steady state. Our goal is to simulate a scenario that isolates the fact that improving one generation's level of skills has intergenerational dynamic effects that accumulate over time. In our model, improving one generation's skills creates higher-skilled and higher-income parents (which may invest more in their children and makes these investments more productive) and higher-income neighbors (increasing neighborhood qualities). In line with this, the effect on labor productivity almost triples (from 0.8 to 2.3 percent, as shown in Column 9) and welfare gains increase by 2.2 percentage points (from 4.4 percent to 6.6 percent), which is about one-third of the "RCT" gains of 5.7 percent.

Next, the fourth and fifth rows show the impact of the two equilibrium forces that are expected to reduce the benefits impacts of voucher programs. The simulation in the fourth row allows for housing price adjustments. Rent in $n = 2$ (i.e., τ_2) increases by 3.5 percent, reducing moves to the advantaged neighborhood relative to the scenario in which rents do not adjust (a 6.8 percent increase rather than 11.1 percent). Taking into account housing prices reduces the welfare gains from voucher programs by 1.3 percentage points (from 6.6 to 5.3). Similarly, the fifth row shows that there is a 2 percentage point decrease in welfare benefits when the simulation allows neighborhood quality to adjust in response to resorting. As more low-income parents and children relocate to $n = 2$, income per capita in the advantaged neighborhood is reduced and neighborhood quality declines by 4.5 percent. Taking into account these resorting effects reduces the average child's neighborhood quality, thereby cutting the labor productivity gains by half (from 2.3 to 1.3 percent, as shown in Column 9). These two rows of results collectively suggest that subsequent household resorting may reduce the potential of voucher programs to generate welfare gains by about half (from 6.6 to 3.3 percent).

The sixth row allows for adjustments in wages and interest rates which are the last general equilibrium forces that we consider. Although the policy leads to a relatively small (0.8 percent) increase in the interest rate and a small decrease in wages (-0.4 percent), this dampens the negative equilibrium effects of reduced neighborhood quality and increased taxation. Thus, welfare gains increase by 1 percentage point (from 3.3 to 4.3), about one-sixth of the "RCT" gains of 5.7 percent.

Finally, we conclude our decomposition analysis by also studying impacts on income inequality and upward mobility. We measure income inequality using the variance of log-lifetime-after-tax-earnings. Upward mobility is defined as the probability that a child born to parents in the bottom quintile of the income distribution is in the top income quintile during the working stage

of their life. Inequality is reduced by 2.7 percent when taking into account all equilibrium forces (last row). While this is smaller than the 7.3 percent reduction in a scenario without equilibrium forces (first row), it is worth noting that 2.7 percent is about half as large as the percent difference in after-tax inequality between Sweden and the U.S. (Table 3 of Krueger et al., 2010). In contrast to this pattern of results, upward mobility effects are similar in the short-run, small-scale version and the large-scale, long-run version. The effect is 29 percent, which is approximately half of the standard deviation in upward mobility across U.S. Census tracts (Chetty et al., 2018). The fact that the largest effects are observed on upward mobility is potentially not surprising given that the voucher policy exposes children of low-income parents to a higher quality neighborhood.

To summarize, a government housing program can lead to large welfare gains in the long-run steady equilibrium. Equilibrium increases in rent and reductions in neighborhood quality have the most important negative impacts. Relative to a scenario without such forces (e.g., an RCT implementation), these forces reduce the gains to vouchers by one-half. Higher taxation also reduces welfare gains by a quarter. However, these negative forces are partially compensated by the fact that a long-run and large scale voucher program enhances child development and creates better parents and neighbors for future generations. This intergenerational dynamic effect, together with the general equilibrium effects on wages and the interest rate, is sufficiently large to offset most countervailing equilibrium forces, making welfare gains only 25 percent smaller than in an RCT implementation.

5.1.2 Voucher Transition Dynamics

This section evaluates the transition dynamics associated with implementing the benchmark housing voucher program described above. A logical concern is that the welfare benefits achieved in the long-run steady state may take too long to accrue. This matters for understanding the political economy issues at play with the housing voucher program that we consider.


To study dynamics, we simulate the model where the government unexpectedly introduces the benchmark housing voucher program with associated (i.e., the one associated with the largest long-run welfare gains). Note that the steady-state change in labor income tax may not be enough to balance the government's budget initially because the pool of skills in the economy takes time to increase. Due to this, the government is assumed to adjust taxes (by adjusting tax parameter λ) every period in order to achieve a balance budget in each period in the transition.

Figure 5 plots the impacts of the voucher program for newborns and future cohorts as a solid (blue) line.⁴⁰ Here, cohort "0" is the first cohort born at the time the policy is introduced, and

⁴⁰The figure also reports effects for an alternative policy, a place-based wage-subsidy, discussed in Section 5.2.

cohort x (for all $x > 0$) refers to the cohort born x periods after the policy is introduced. Panel (a) shows that welfare increases by about 4.5 percent for cohort 0. These gains initially decline for subsequent cohorts to around 3 percent before rebounding. The impact on welfare stabilizes to the steady state increase of 4.3 percent for the twentieth cohort that is born after the introduction of the housing voucher program. Panel (b) shows that productivity increases cohort-by-cohort until the steady-state level is achieved. The first cohort's productivity increases by about 0.5 percent. Productivity is then almost unchanged until a jump (to about 1 percent) is observed for the first cohort born to the parents who received the intervention (i.e., those born 28 years after the policy is introduced). A second jump (to almost 1.2 percent) takes place afterwards for the first cohort born that had grandparents who received the voucher subsidy. These jumps in the productivity demonstrate the mechanism driving long-run intergenerational gains: exposing a child to a better neighborhood today creates a better parent for the next generation.

Figure 5: Transition Analysis Comparison: Effects of the Benchmark Policies by Cohort



Notes: This figure presents an analysis of the transition dynamics associated with the two government interventions. Results for the rent voucher program that generates the highest steady-state equilibrium welfare gains (detailed in Section 5.1) are reported in solid (blue). Results for the place-based wage subsidy program that generates the highest steady-state equilibrium welfare gains (detailed in Section 5.2) are reported in dashed (red). All effects of the programs are at the cohort level starting with the first cohort born at the introduction of the policy (i.e., cohort 0) up to the fiftieth cohort born afterward. The x -axis on the figure indicates the relevant cohort. These effects represented by the y -axis are calculated as the difference in a given outcome compared to the steady-state where no housing voucher program existed. Note that welfare is measured as consumption equivalence (see Section C for further details).

Why do welfare gains initially decline before rebounding? Figure 6 provides an explanation by showing the dynamics of several key economic outcomes by time period. Here, period “0” is the time when the rent voucher policy is introduced. There are two key points from this analysis. First, Panels (a) and (b) show that neighborhood quality decreases and rent increase shortly after the

policy is introduced which lowers welfare gains.⁴¹ Theoretically, these effects on neighborhood conditions and rent stem from moving costs that slow resorting. This prediction is confirmed in Panel (d) which shows gradual change in population shares in $n = 2$.⁴² The second key finding can be observed from the transition dynamics for the capital stock and GDP illustrated in Panels (e) and (f), respectively. Capital in the disadvantaged neighborhood partially recovers over time which drives a recovery in GDP. Yet, these changes lag the resorting that drives changes in rents and quality. To summarize, the initial decline in welfare gains is due to decreases in neighborhood quality and increases in rent that are eventually offset through capital stock and GDP growth.

5.2 Evaluating Place-Based Policies

The goal of this section is to study the effects of place-based policies in general equilibrium. The policy that we study is a wage subsidy, \tilde{w}_s , for individuals who live (and work) in the disadvantaged area. Specifically, the policy is such that in equilibrium workers in $n = 1$ earn $w_1 = (1 + \tilde{w}_s)w_2$. As in our analysis of housing vouchers, the government finances this program by adjusting labor taxes (λ) to ensure the budget remains at its initial steady state level.


In line with our prior analysis, we aim to identify the wage subsidy program that generates the highest steady state welfare (under the veil of ignorance). We determine this by varying the subsidy rate from 1 to 20 percent (in intervals of 1 percent) and simulating the equilibrium in each case. For each simulation, we calculate percent changes for various equilibrium outcomes relative to the baseline scenario in which there is no wage subsidy program.

Figure 7 reports results from the various wage subsidies that we consider. The vertical dashed (black) line in Panel (a) shows that the highest steady state welfare gain is 0.5 percent when the wage subsidy is set to 11 percent. The remaining panels on the top row show there are important impacts on residential sorting. At this welfare maximizing subsidy rate, Panel (b) shows that neighborhood quality increases in the disadvantage area by 17.5 percent. As shown in Panel (c), this is driven by relocation of relatively higher-skilled workers who are induced to move from the advantaged area. Note that Panel (c) also shows that the share of children in the advantaged also decreases as their parents are drawn by the higher wages, lower rents and the newly realized increases in neighborhood quality.

⁴¹In contrast, Panel (c) shows that equilibrium taxes increase immediately and do not change substantively over time (i.e., they are relatively constant at a 16 percent increase). This rules out that taxes play a role in explaining the initial decrease in welfare gains observed in Figure 5.

⁴²Panel (d) also shows population shares by type of household. Families with children (particularly children who are young) are the first to move. Population shares for older individuals increase much more slowly. This pattern for the elderly is driven by long-run dynamics: young individuals with children initially move into $n = 2$, and these cohorts subsequently stay in this location during old age (potentially due to moving costs).

Figure 6: Transition Analysis: Effects of the Benchmark Voucher Program by Period




Notes: This figure presents an analysis of the transition dynamics associated with the housing voucher program that generates the highest steady-state equilibrium welfare gains. We report effects of the voucher program at the period level starting with the first period when the policy is in effect (i.e., period 0) up to the fiftieth period afterward. The x -axis on the figure indicates the period of interest. The effects represented by the y -axis are calculated as the difference in a given outcome compared to the initial steady state (where no housing voucher program existed). Note that the results in Panel (d) show population shares in the advantaged neighborhood ($n = 2$) for the following groups: all individuals, young individuals with no children (i.e., $j = 5 - 7$), young individuals with children ($j = 8 - 11$), and old working-age individuals without children ($j = 12 - 16$).

The results in Figure 7 also show how welfare and other equilibrium outcomes change as a function of the subsidy rate. Panel (a) shows that the steady state welfare gains increase at a relatively constant rate before decreasing when the subsidy exceeds 11 percent. The results in the bottom row of panels explores forces that could drive the welfare gains. Labor productivity plays little role as Panel (d) shows that productivity gains are consistently small under various levels of subsidies. Rather, the main source of welfare gains is likely due to reduced inequality.⁴³ Panel (e) shows that the variance of log-lifetime-earnings is reduced by approximately 9 percent when the wage subsidy is set to 11 percent. Although further reductions in inequality are possible at higher subsidy levels, the results also illustrate the importance of taxes. Panel (f) shows that the average marginal labor tax rate increases steadily in order to raise revenue for the wage

⁴³Given that low-income individuals tend to live in $n = 1$, this effect may also be interpreted as insurance against negative shocks that can reduce income.

subsidy program. The negative effect of increased tax distortions seems to dominate the gains from inequality reductions at relatively high level of taxes.

Figure 7: Long-Run Effects of Alternatives Wage Subsidy Policies



Notes: This figure presents results from an analysis of the effects of various wage subsidy policies on equilibrium outcomes. All effects are calculated by comparing the difference in outcomes between a simulation for a given wage subsidy level and the baseline scenario where there is no government intervention. The x-axis in each panel shows the level of the wage subsidy for each simulation. We vary the subsidy rate from 2 to 24 percent (in intervals of 2 percent). The vertical dashed (black) line in each panel indicates the subsidy rate (12 percent) that achieves the highest steady state welfare gain.

5.2.1 Wage Subsidy Decomposition

Next, we turn to analyzing the equilibrium forces that drive the highest steady state welfare gains achieved with a wage subsidy policy. Our analysis mirrors the approach we use to study housing vouchers. The 11 percent wage subsidy that generates the highest steady state welfare of 0.5 percent is defined as the benchmark and we compare this to several simulations that shut down other channels in the model.

Table 6 presents the wage subsidy policy decomposition results. The first row reports a simulation where the 11 percent wage subsidy is introduced for one generation (starting from the initial steady-state without a wage subsidy program), there are no general equilibrium effects (e.g., neighborhood quality and prices do not change), and the government does not need to bal-

ance its budget. Effects are evaluated for the children of the single generation that is affected by the wage subsidy. In this scenario, children are affected by the wage subsidy due to changes in their parents behavior such as new patterns of investment or neighborhood relocation. As in our previous analysis, the next rows report economic outcomes allowing for additional equilibrium channels, with the final bold row providing results with all general equilibrium effects.

Table 6: Wage Subsidy: Decomposition of Equilibrium Forces for Main Economic Outcomes

Equilibrium Forces					Change from Initial Steady State (%)						
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Long Run	Taxes	Real Estate	Neigh. Qual.	Prices	Share in $n = 2$	Neigh. Qual, $n = 1$	Income Ineq.	Upward Mobility	Labor Prod.	Welf. Adults	Welf. Children
X	X	X	X	X	-3.3	0.0	1.6*	-4.1*	-1.5*	0.7	-1.4*
X	X	X	✓	X	-4.0	6.0	0.4*	6.5*	-0.4*	0.8	-0.5*
X	X	✓	✓	X	-2.2	6.9	-1.0*	12.6*	0.1*	0.7	0.0*
X	✓	✓	✓	X	-2.2	7.0	-1.0*	12.4*	0.1*	0.2	-0.2*
✓	✓	✓	✓	X	-6.4	17.5	-8.7	21.6	0.0	-	0.5
✓	✓	✓	✓	✓	-6.4	17.5	-8.7	21.6	0.0	-	0.5

Notes: This table presents results that decompose how various equilibrium forces affect economic outcomes under the wage subsidy program that generates the highest steady-state welfare gains. Each row provides results from a separate simulation of the model. Columns 1-5 describe which equilibrium force is shut down in each simulation. Columns 6-12 report changes (in percent) in outcomes calculated by comparing each simulation to the initial steady state (where there is no housing voucher program). When calculating short-run effects, GDP is calculated as lifetime earnings for the children of the single generation that is affected by the policy. In such cases, we do not report effects on the capital stock as these depend on selecting a specific time period for measurement. The asterisks highlight that short-run effects are evaluated for children of the cohort that received the policy intervention.

The main finding from our decomposition is that most of the equilibrium forces that we study have a role in determining the 0.5 percent highest steady state welfare gain. In the top row, Column 12 shows that the impact on welfare stands at negative 1.4 percent when the program is implement for a single generation without equilibrium effects. Intuitively, this occurs because the share of children living in the more advantaged neighborhood $n = 2$ decreases which reduces the later-life labor productivity of this single generation whose parents are eligible for the wage subsidy. The next row shows that this impact falls in magnitude to a 0.5 percent reduction when the simulation allows for neighborhood quality to adjust. In this scenario, Columns 7 shows that neighborhood quality in the low-income neighborhood increases by 6 percent. Similarly, the third row shows that allowing for changes in rental prices in the housing market further reduces the negative impacts because fewer children move away from the advantaged neighborhood (due to the price increases in $n = 1$ in this scenario). The fourth row shows that the need to increase taxes to finance the subsidy reduces welfare gains by 0.2 percent.

The fifth row shows that implementing the program for the long-run shifts the welfare gains

from negative to positive (from -0.2 to 0.5 percent, as shown in Column 12). Two factors drive this increase in welfare in this scenario. One is that children are able to benefit from the subsidy directly (rather than solely due to changes in parental behavior) when they reach adulthood. Another is that there are larger reductions in inequality. As seen in the first few rows, we find that a short-run program has little impact on inequality even when we allow for equilibrium effects on neighborhood quality or housing markets. In contrast, instituting the wage policy in the long-run is what drives most of the impact on inequality and over half of the impact on upward mobility. Finally, the last row shows that accounting for wage and interest rate changes does not seem important for this policy.


5.2.2 Wage Subsidy Transition Dynamics

In the next component of our analysis, we study the transition dynamics associated with implementing the 11 percent benchmark wage subsidy policy in general equilibrium. To begin this discussion, we return to Figure 5 which also plots the impact of the wage subsidy for newborns and future cohorts as a dashed (red) line. Panel (a) shows that welfare gains fluctuate around 0 for cohorts born soon after the policy is implemented, and they grow slowly until they stabilize at the long run gains of 0.5. In line with this, Panel (b) shows that the initial cohorts experience productivity losses which are attenuated for later-born cohorts.⁴⁴

The relatively constant gains in welfare and productivity by cohort are in line with immediate and relatively rapid adjustments in key economic outcomes by period reported in Figure 8. Panel (a) shows that taxes jump when the policy is put in place and increase until stabilizing at 2.5 percent. The effects on neighborhood quality in Panel (b) occur slightly more gradually. After 10 periods, the impact on neighborhood quality is 21 percent, decreasing afterwards and stabilizing around 18 percent. Panel (c) shows that rent increases over time, until it reaches an increase of about 30 percent. Panel (d) demonstrates that resorting drives the impact on neighborhood quality (and rent increases). The share of young workers declines almost immediately in the advantaged $n = 2$ neighborhood as these workers seek to fully exploit the benefits of the wage subsidy. Relocation takes slightly longer (about 10 periods) for most of the remaining age groups.

⁴⁴Note that the cohort-level effects initially fluctuate due to moving costs. Just after the policy is introduced, older individuals and families with children do not move to $n = 1$. Instead, the young without children move and are more likely to stay in $n = 1$ even in old age. This pattern increases living costs and reduces neighborhood quality. At this point, the next cohorts are less likely to move to $n = 1$, and productivity impacts of the policy move in the opposite direction. This process repeats until population shares in each neighborhood converge to the long-run equilibrium.

Figure 8: Transition Analysis: Effects of the Benchmark Wage Subsidy Policy by Period



Notes: This figure presents an analysis of the transition dynamics associated with the place-based wage subsidy program that generates the highest steady-state equilibrium welfare gains. We report effects of the wage subsidy program at the period level starting with the first period when the policy is in effect (i.e., period 0) up to the fiftieth period afterward. The x -axis on the figure indicates the period of interest. The effects, represented by the y -axis are calculated as the difference in a given outcome compared to the initial steady state (where no housing voucher program existed). Note that the results in Panel (d) show population shares in the advantaged neighborhood ($n = 2$) for the following groups: all individuals, young individuals with no children (i.e., $j = 5 - 7$), young individuals with children ($j = 8 - 11$), and old working-age individuals without children ($j = 12 - 16$).

5.3 Robustness Exercises: Welfare Analysis and Parameter Sensitivity

The results so far show that rent vouchers and wage subsidies can lead to long-run welfare gains. Moreover, we find that the gains from rent vouchers are nearly nine times as large as those from wage subsidies. This section studies the robustness of our welfare results when our calibrated model relies on alternative parameter values.

We begin by evaluating the change in welfare gains when we increase each parameter in the model parameters by one percent while holding other parameters constant. This approach to parameter sensitivity analysis follows [Andrews et al. \(2017\)](#) and [Elenev et al. \(2020\)](#). Appendix Table A3 reports our estimates of welfare changes with separate rows for the the benchmark rent voucher and subsidy programs. One main finding is that changes in most parameters have the same signed effects on welfare for both policies. Moreover, the effects on each policy are often small and similar in magnitude. Together, these results suggests that most changes of individual

parameter values will not cause us to reach different conclusion as to whether a rental subsidy policy generates higher steady-state welfare.

While the sensitivity results provide reassurance for most parameters, it is worth discussing potential concern over the results for two parameters: altruism and the housing supply elasticity. These parameters are among the few that have opposite signed effects on welfare. Moreover, the housing supply elasticity is of particular concern given that this is a key parameter that governs the negative impacts of rental vouchers. That is, large reductions in this parameter have distinctly negative impacts on welfare under the rent subsidy program because there are increased displacement effects when it is harder to accommodate growth.

First, in the case of altruism, a back-of-the-envelope calculation is informative as to whether large changes in this parameter are empirically reasonable. Specifically, $\tilde{\beta}$ would need to more than double to cause the welfare gains of the rent voucher program to fall below the wage subsidy. While a formal analysis would require knowing the standard deviation of altruism, we can assess the plausibility of this change by evaluating how such an increase would affect non-welfare moments and compare the resulting (changed) moments to empirical benchmarks. For example, increasing $\tilde{\beta}$ by one percent would increase the parental-transfer estimation moment (i.e., the average parental transfer as a share of average income) by 2.8 percent (from 128.5 percent in our benchmark voucher program). Thus, doubling $\tilde{\beta}$ would require increasing this parent-transfer moment to over 400 percent—relatively far from the empirical benchmark of 125.4 percent.

Second, we provide additional sensitivity analysis to gauge how our conclusions are shaped by substantial changes in the housing supply elasticity. This exercise is motivated in part by recent research that has found evidence of notably smaller elasticities and documented housing supply responses that vary substantially across large metro areas of the U.S. (Baum-Snow and Han, 2021). Table 7 reports welfare estimates in scenarios where housing supply responses are more constrained. For comparison, the first row begins by reproducing the main welfare estimates using the average elasticity from Saiz (2010). The second row also shows that welfare is only slightly smaller in the scenario where the supply elasticity is decreased by just one percent (as in Appendix Table A3). The next rows use a more conservative estimate of housing supply elasticity based on analysis from Baum-Snow and Han (2021). Using more recent data, these authors estimate an elasticity of 0.30—i.e., 80 percent smaller than the estimate of 1.75 in our main analysis.⁴⁵ Note that the third and fourth rows provide results where we also re-estimate the

⁴⁵Baum-Snow and Han (2021) argue that the main explanation for the difference in their result is that their analysis estimates a short-run elasticity using data from 2000-2010 (whereas Saiz (2010) estimate an elasticity using data from 1970-2000). Another distinction is that Baum-Snow and Han (2021) conduct their analysis at the Census tract-level analysis rather than the MSA-level as in Saiz (2010). They suggest the level of geography for their analysis

internally calibrated parameters to match the same moments as in Table 2.

Table 7: Parameter Sensitivity Analysis for the Housing Supply Elasticity

	(1)	(2) LR Welfare Gains	
		Elasticity	Rent Voucher
Baseline (Saiz, 2010)	1.75	4.35	0.49
1% Lower Supply Elasticity	1.73	4.33	0.51
Baum-Snow and Han (2021) (Re-estimated)			
Same policies	0.30	2.9	1.6
New optimal policies	0.30	2.9	2.6

Notes: This table reports steady-state welfare estimates for rent voucher and wage subsidy programs. The first three rows report welfare rely on different assumed housing supply elasticities but the policies have the program rules (e.g., eligibility for vouchers or the amount of the wage subsidy) that we specify in our main analysis in Section 5. The fourth row estimates welfare when the housing supply elasticity is lower than the parameter value in our main analysis and the voucher and wage subsidy program rules are set to the values that maximize steady state welfare.

The main finding from our analysis using a notably small housing supply elasticity is that welfare gains associated with vouchers decrease notably. The third row of Table 7 shows that the welfare gains fall from 4.35 to 2.9 percent for the voucher program. Intuitively, this occurs because a lower supply elasticity leads to larger housing price increases, which makes the policy more expensive and forces more households to move out of the advantaged neighborhood. We also find that the gains of wage subsidies increase from 0.49 to 1.6 percent. In this case, a lower elasticity leads to larger reduction in housing prices in the advantaged neighborhood (as people choose to move out), which increases the welfare gains of the place-based policy. The fourth row extends on these results by reporting welfare after we solve for new optimal (highest steady state welfare) policies in the world where the housing elasticity is set to 0.30. While rental voucher gains remain unchanged (i.e., the optimal policy remains the same), the welfare benefits of wage subsidies increase further to 2.6 percent (since the optimal policy is now one with higher wage subsidies).

Overall, we reach the following conclusions from the sensitivity analysis in this section. For nearly all the parameters, our conclusions are robust to empirically justified changes. Our sensitivity analysis shows that changes in the parameters generally have small and same-signed impacts on the welfare gains associated the housing programs we study. The notable exception to this pattern of results is the housing supply elasticity. Using a more conservative estimate of

is not the main driver of the difference in their estimated elasticity.

the housing supply elasticity reshapes the relative merits of vouchers versus place-based investments. When we depart from the standard estimated elasticity from [Saiz \(2010\)](#), the two programs generate comparable welfare gains using the smaller housing supply elasticity from ([Baum-Snow and Han, 2021](#)). Moreover, given that housing supply elasticities vary across regions (e.g., at the low-end of the distribution, [Baum-Snow and Han \(2021\)](#) estimate an elasticity of 0.12 for both San Francisco and New York), our sensitivity results suggest that the balance may tip in favor of wage subsidies in the most highly constrained U.S. cities.

5.4 Discussion: Comparing Voucher and Wage Subsidy Policies

Why is it possible to achieve higher welfare gains with a voucher program relative to a place-based wage subsidy? The analysis so far shows that the highest steady state welfare for a housing voucher program is 4.3 percent larger relative to our baseline scenario where the government does not intervene in the housing market. In contrast, we find that the highest steady state improvement with a wage subsidy policy is 0.5 percent.


The main explanation is that the highest steady-state-welfare voucher program has larger impacts on labor productivity. Specifically, the voucher program can generate a 1.4 percentage point increase in productivity ([Table 5, Column 10](#)), whereas the wage subsidy does not generate any productivity increase ([Table 6, Column 10](#)). This difference in labor productivity is due to differences in each program's impact on neighborhood conditions (which thereby affects child development). Equilibrium neighborhood quality for the average child increases by 2.2 percent for the voucher program, but it decreases by 0.4 percent for the wage subsidy program.

Notably, the results do not indicate that the relatively larger welfare gains from the voucher program are due to the programs impact on inequality. The benchmark voucher program generates reduces income inequality by 2.7 percent ([Table 5, Column 8](#)). In contrast, the wage subsidy has a larger magnitude impact, decreasing inequality by 8.7 percent ([Table 6, Column 8](#)).

Given the welfare gains associated with both programs, a final exercise that we undertake sheds further light on the political economy associated with these policies. A natural concern is that the policies may have heterogeneous effects on welfare for adults alive at the introduction of the policy (hereafter "incumbent adults"). Heterogeneity in the gains from either program implies that policymakers may face tradeoffs when considering whether to implement voucher or place-based wage subsidies in a democratic system.

[Figure 9](#) reports welfare gains from voucher and place-based subsidies for incumbent adults. The results in the top row highlight important heterogeneity in the effects of both programs. For vouchers, the results in [Panel \(a\)](#) show that welfare gains are concentrated among the individuals

Figure 9: The Welfare Effects of Rent and Wage Subsidy Policies on Incumbent Cohorts



Notes: This figure presents an analysis of welfare gains for adults alive at the time the highest-welfare voucher program is introduced (i.e., incumbent adults). Panels (a) and (b) plot the effects by age and by age and neighborhood of residence, respectively. In both panels, the x-axis indicates the age of an adult at the policy’s introduction.

who will soon have children (i.e., those aged 16–27) and those who already have children (i.e., those aged 28–43). Older individuals, instead, tend to lose from the new housing voucher policy.⁴⁶ Panel (b) shows that the wage subsidy has a different pattern of impacts. Notably, the average benefit for each cohort alive at the policy’s introduction is always positive (although the gains are larger for younger cohorts).

The second row extends these results by showing welfare impacts by age *and* the initial loca-

⁴⁶This result is *not* driven by failing to allow the parental value function to account for changes to child utility after they have leave the households. To avoid this typical issue when calculating welfare gains in OLG models over the transition, we extend the parents value function to include its effect on their children (and future descendants) when calculating welfare gains for incumbent adults whose children have already become independent (i.e., $j \geq 12$).

tion where individuals reside. The results in Panel (c) show that, within the population of younger individuals, it is those originally living in the disadvantaged neighborhood ($n = 1$) who have larger welfare benefits from the voucher program. The results in Panel (d) show the disadvantaged residents again are the largest beneficiaries of the place-based wage subsidy program. Notably, these results also show that most age cohorts in the advantaged neighborhood have average welfare losses from the wage subsidy program.

What do we conclude from this analysis of incumbent adults? The benchmark voucher program leads to larger long-run welfare gains relative to a place-based wage subsidy, but there is heterogeneity that has political economy considerations. If individuals calculate welfare gains as in our model, a majority of adults would vote against both policies: 56 and 74 percent of incumbent adults are made worse off under the voucher and wage subsidy policies, respectively. This is true for the place-based policy because a majority (approximately 90 percent) of individuals live in $n = 2$ and, as noted above, suffer welfare losses. While vouchers have more political support (44 percent) from a democratic voting perspective, it is worth noting that this program may be controversial in terms of welfare impacts. On average, incumbents who are made worse off under vouchers have welfare losses of 2.5 percent. This is about 10 times as large as the average loss of 0.24 percent for incumbents who fail to benefit from the wage subsidy. Overall, these results suggest that there are important tradeoffs associated with each program. Further, either policy may only be acceptable to a majority of individuals in a setting where the government is able to borrow to initially finance either program and increase taxation in the future—a finding which is similar to the conclusions regarding early childhood education programs (Daruich, 2020).

6 Conclusion

This paper provides a new quantitative assessment of policies that aim to shape neighborhood quality for children. Building on prior theoretical research that studies inequality and neighborhoods (Benabou, 1996a; Durlauf, 1996; Fernandez and Rogerson, 1996), our analysis focuses on a spatial equilibrium model that features overlapping generations and incorporates endogenous childhood development. We calibrate the model based on U.S. data and use simulations to study the long-run and large-scale impacts housing vouchers and location-specific wage subsidies, two types of government policies common around the world. The programs that we study represent distinct antipoverty strategies: a people-based approach that provides assistance directly to low-income families and a place-based approach that targets government resources at a local area.

Our core finding is that housing vouchers and place-based wage subsidies can increase welfare in the long-run despite several potentially countervailing equilibrium forces. These welfare gains

occur because both programs increase the average neighborhood quality for children, thereby creating better parents and neighbors for future generations. We find that housing vouchers can generate larger gains in welfare relative to what is feasible with a wage-subsidy approach.

It is worth highlighting two important caveats for our main conclusions. First, our sensitivity analysis suggests that the long-run gains of housing voucher programs can be notably smaller in cities with limited housing supply responses. We provide suggestive evidence that wage subsidy programs may generate higher welfare in constrained cities such as San Francisco or New York. Second, although we find that housing vouchers represent a particularly promising long-run approach to increasing welfare for the average city, our analysis of transition dynamics suggests there may be limited political support for this type of program. For adults alive at the introduction of the benchmark voucher program that we study, young cohorts achieve welfare gains while older cohorts are worse off. We similarly find limited democratic support for a place-based wage subsidy program. These political economy results suggest that policymakers face important tradeoffs when enacting people- or place-based government interventions.

References

- Abbott, Brant.** (2021). ‘Incomplete markets and parental investments in children’, *Review of Economic Dynamics* .
- Abbott, Brant, Gallipoli, Giovanni, Meghir, Costas and Violante, Giovanni L.** (2019). ‘Education policy and intergenerational transfers in equilibrium’, *Journal of Political Economy* 127(6), 2569–2624.
- Agostinelli, Francesco, Doepke, Matthias, Sorrenti, Giuseppe and Zilibotti, Fabrizio.** (2020), It takes a village: the economics of parenting with neighborhood and peer effects, Technical report, National Bureau of Economic Research.
- Aliprantis, Dionissi and Carroll, Daniel R.** (2018). ‘Neighborhood dynamics and the distribution of opportunity: Neighborhood dynamics’, *Quantitative Economics* 9(1), 247–303.
- Andrews, Isaiah, Gentzkow, Matthew and Shapiro, Jesse M.** (2017). ‘Measuring the sensitivity of parameter estimates to estimation moments’, *The Quarterly Journal of Economics* 132(4), 1553–1592.
- Baland, Jean-Marie and Robinson, James.** (2000). ‘Is Child Labor Inefficient?’, *Journal of Political Economy* 108(4), 663–679.
- Barro, Robert J. and Becker, Gary S.** (1989). ‘Fertility Choice in a Model of Economic Growth’, *Econometrica* 57(2), 481–501. Publisher: [Wiley, Econometric Society].
- Bartik, Timothy J.** (2020). ‘Using Place-Based Jobs Policies to Help Distressed Communities’, *Journal of Economic Perspectives* 34(3), 99–127.
- Baum-Snow, Nathaniel and Han, Lu.** (2021). ‘The microgeography of housing supply’, *Work in progress* .
- Benabou, Roland.** (1996a). ‘Equity and Efficiency in Human Capital Investment: The Local

- Connection', *The Review of Economic Studies* 63(2), 237.
- Benabou, Roland.** (1996*b*). 'Heterogeneity, Stratification, and Growth: Macroeconomic Implications of Community Structure and School Finance', *The American Economic Review* 86(3), 584–609.
- Biasi, Barbara.** (2019), School Finance Equalization Increases Intergenerational Mobility: Evidence from a Simulated-Instruments Approach, Technical Report w25600, National Bureau of Economic Research, Cambridge, MA.
- Boustan, Leah Platt.** (2010). 'Was Postwar Suburbanization "White Flight"? Evidence from the Black Migration', *The Quarterly Journal of Economics* 125(1), 417–443.
- Busso, Matias, Gregory, Jesse and Kline, Patrick.** (2013). 'Assessing the Incidence and Efficiency of a Prominent Place Based Policy', *American Economic Review* 103(2), 897–947.
- Caucutt, Elizabeth M, Lochner, Lance, Mullins, Joseph and Park, Youngmin.** (2020), Child skill production: Accounting for parental and market-based time and goods investments, NBER Working Paper.
- Chetty, Raj, Friedman, John N, Hendren, Nathaniel, Jones, Maggie R and Porter, Sonya R.** (2018), The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility, Working Paper 25147, National Bureau of Economic Research.
- Chetty, Raj and Hendren, Nathaniel.** (2018*a*). 'The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects*', *The Quarterly Journal of Economics* .
- Chetty, Raj and Hendren, Nathaniel.** (2018*b*). 'The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates', *The Quarterly Journal of Economics* 133(3), 1163–1228.
- Chetty, Raj, Hendren, Nathaniel and Katz, Lawrence F.** (2016). 'The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment', *American Economic Review* 106(4), 855–902.
- Chyn, Eric.** (2018). 'Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children', *The American Economic Review* 108(10), 3028–3056.
- Chyn, Eric, Hyman, Joshua and Kapustin, Max.** (2019). 'Housing Voucher Take-Up and Labor Market Impacts', *Journal of Policy Analysis and Management* 38(1), 65–98.
- Collinson, Robert, Ellen, Ingrid Gould and Ludwig, Jens.** (2015), Low-Income Housing Policy, Working Paper 21071, NBER.
- Cunha, Flavio.** (2013), Investments in Children When Markets Are Incomplete.
- Cunha, Flavio, Heckman, James J. and Schennach, Susanne M.** (2010). 'Estimating the Technology of Cognitive and Noncognitive Skill Formation', *Econometrica* 78(3), 883–931.
- Daruich, Diego.** (2020), The Macroeconomic Consequences of Early Childhood Development Policies.
- Daruich, Diego and Fernández, Raquel.** (2020), Universal basic income: A dynamic assessment, Technical report, National Bureau of Economic Research.
- Davis, Morris A, Gregory, Jesse and Hartley, Daniel A.** (2019), The Long-Run Effects of Low-Income Housing on Neighborhood Composition.
- Davis, Morris, Gregory, Jesse, Hartley, Daniel and Tan, Kegan T.K.** (2021), Neighborhood Effects and Housing Vouchers, Technical Report w28508, National Bureau of Economic Research.

search, Cambridge, MA.


- Derenoncourt, Ellora.** (2019). ‘Can you move to opportunity? Evidence from the Great Migration’.
- Diamond, Rebecca and McQuade, Tim.** (2019). ‘Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development’, *Journal of political economy* p. 55.
- Durlauf, Steven N.** (1996). ‘A Theory of Persistent Income Inequality’, *Journal of Economic Growth* 1(1), 75–93. Publisher: Springer.
- Eckert, Fabian and Kleineberg, Tatjana.** (2021), Education Policies in Spatial General Equilibrium.
- Elenev, Vadim, Landvoigt, Tim and Van Nieuwerburgh, Stijn.** (2020), A macroeconomic model with financially constrained producers and intermediaries, Technical report, National Bureau of Economic Research.
- Fernandez, Raquel and Rogerson, Richard.** (1996). ‘Income Distribution, Communities, and the Quality of Public Education’, *The Quarterly Journal of Economics* 111(1), 135–164. Publisher: Oxford University Press.
- Fernandez, Raquel and Rogerson, Richard.** (1998). ‘Public Education and Income Distribution: A Dynamic Quantitative Evaluation of Education-Finance Reform’, *The American Economic Review* 88(4), 813–833. Publisher: American Economic Association.
- Fogli, Alessandra and Guerrieri, Veronica.** (2019), The End of the American Dream? Inequality and Segregation in US Cities, Working Paper 26143, National Bureau of Economic Research.
- Heathcote, Jonathan, Storesletten, Kjetil and Violante, Giovanni L.** (2017). ‘Optimal tax progressivity: An analytical framework’, *The Quarterly Journal of Economics* 132(4), 1693–1754.
- Howell, Penny L. and Miller, Barbara B.** (1997). ‘Sources of Funding for Schools’, *The Future of Children* 7(3), 39–50. Publisher: Princeton University.
- Hoxby, Caroline M.** (2001). ‘All School Finance Equalizations are Not Created Equal*’, *The Quarterly Journal of Economics* 116(4), 1189–1231.
- Katz, Lawrence F., Kling, Jeffrey R. and Liebman, Jeffrey B.** (2001). ‘Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment’, *The Quarterly Journal of Economics* 116(2), 607–654.
- Kling, Jeffrey R., Liebman, Jeffrey B. and Katz, Lawrence F.** (2007). ‘Experimental Analysis of Neighborhood Effects’, *Econometrica* 75(1), 83–119.
- Kling, Jeffrey R., Ludwig, Jens and Katz, Lawrence F.** (2005). ‘Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment’, *The Quarterly Journal of Economics* 120(1), 87–130.
- Krueger, Dirk, Perri, Fabrizio, Pistaferri, Luigi and Violante, Giovanni L.** (2010). ‘Cross-sectional facts for macroeconomists’, *Review of Economic dynamics* 13(1), 1–14.
- Loury, Glenn C.** (1981). ‘Intergenerational Transfers and the Distribution of Earnings’, *Econometrica* 49(4), 843–867.
- Ludwig, Jens, Duncan, Greg J., Gennetian, Lisa A., Katz, Lawrence F., Kessler, Ronald C., Kling, Jeffrey R. and Sanbonmatsu, Lisa.** (2013). ‘Long-Term Neighborhood Effects on Low-Income Families: Evidence from Moving to Opportunity’, *NBER*.

- McDaniel, Cara.** (2007), Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003, Technical report.
- Neumark, David and Simpson, Helen.** (2015), Place-Based Policies, in 'Handbook of Regional and Urban Economics', Vol. 5, Elsevier, pp. 1197–1287.
- Price, Joseph.** (2008). 'Parent-Child Quality Time Does Birth Order Matter?', *Journal of Human Resources* 43(1), 240–265. Publisher: University of Wisconsin Press.
- Rothenberg, Thomas J.** (1971). 'Identification in Parametric Models', *Econometrica* 39(3), 577–591. Publisher: [Wiley, The Econometric Society].
- Saiz, Albert.** (2010). 'The geographic determinants of housing supply', *The Quarterly Journal of Economics* 125(3), 1253–1296.
- Salvi, Angelica, Adema, Willem, Frey, Valerie and Ferraro, Valeria.** (2016), Policies to promote access to good-quality affordable housing in OECD countries, Technical Report 176, OECD, Paris.
- Sampson, Robert J. and Groves, W. Byron.** (1989). 'Community Structure and Crime: Testing Social-Disorganization Theory', *American Journal of Sociology* 94(4), 774–802.
- Sampson, Robert J., Morenoff, Jeffrey D. and Gannon-Rowley, Thomas.** (2002). 'Assessing "Neighborhood Effects": Social Processes and New Directions in Research', *Annual Review of Sociology* 28(1), 443–478.
- Wilson, William J.** (1987), *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy*, University of Chicago Press.
- Zheng, Angela and Graham, James.** (2020). 'Public Education Inequality and Intergenerational Mobility', *American Economic Journal: Macroeconomics* .

Online Appendix

A Appendix Figures

Figure A1: Economic Outcomes Across Neighborhoods in Chicago



Notes: Panel (a) plots median household income from the 1990 U.S. Decennial Census. Panel (b) plots estimates of mean household income ranks for children who grew up in the tract and had parents with household income at the 25th percentile of the national income distribution. This measure of “upward mobility” for children comes from the Opportunity Atlas (Chetty et al., 2018). The measure is specific to children who were born in the 1978-83 cohorts.

A Appendix Tables

Table A1: Estimates of Wage Parameters

	(1)	(2)	(3)
Age	0.0356*** (0.003)		
Age ²	-0.000*** (0.000)		
Inv. Mills Ratio	-1.611*** (0.039)		
Υ		0.999*** (0.021)	
ρ			0.959*** (0.000)
σ_z			0.037*** (0.000)
σ_{η_0}			0.042*** (0.000)
R^2	0.116	0.146	–
# of households	3,052	2,509	2,509
Observations (N)	21,204	19,603	19,603

Notes: This table reports estimates of the parameters of the wage process in our model. Column 1 reports results for the age profile parameters. This is obtained using a sample constructed from the PSID (1968–2016) and regressing wages on age, age-squared, and controls for selection into work based on the Inverse Mills Ratio obtained from a Heckman-selection correction approach. The selection estimator is based on estimating an employment participation equation using the number of children and year-region fixed effects. Column 2 reports estimates of the return to skills. This is obtained using a sample from the NLSY and regressing of the idiosyncratic component of labor productivity ψ_j (measured as a residual based on the age profile estimates from Column 1) on the log of cognitive skills as measured by the AFQT score. Column 3 reports estimates of the parameters that govern the AR(1) process that we assume determines the shock η_j which is the idiosyncratic component of labor productivity. These estimates are obtained from the Minimum Distance Estimator developed by [Rothenberg \(1971\)](#). Standard errors are reported in parentheses. Statistical significance is denoted by: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A2: Estimates of the Elasticity of Time with Respect to Family Income

	(1)	(2)	(3)	(4)
Log(Family Income)	0.105*** (0.006)	0.095*** (0.006)	0.103*** (0.007)	0.097*** (0.007)
R^2	0.154	0.156	0.184	0.190
Observations (N)	26,408	26,408	26,408	26,408
Controls?				
Number and age of children	✓	✗	✗	✗
Household type	✓	✓	✗	✗
Respondent occupation	✓	✓	✓	✗
Commuting zone	✓	✓	✓	✓

Notes: This table reports estimates of the elasticity of quality time spent with children and family income. Columns report estimates that vary based on the set of controls included in the specification. The sample used for the analysis is based on the ATUS, as described in Section 4.1. Note that family income is defined as the combined income of all family members during the last 12 months (including money from jobs, net business income, farm or rent, pensions, dividends, interest, Social Security payments, and any other money received by family members who are 15 years of age or older). Standard errors are reported in parentheses. Statistical significance is denoted by: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A3: Welfare Changes and Parameter Sensitivity Analysis

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		Change in Welfare Given 1 Pct. Increase in Parameter							
Panel A. Parameters:									
	Baseline Welfare Gain	μ	$\tilde{\beta}$	\bar{v}_1	$\bar{\kappa}$	σ_κ	ξ	γ	$\bar{A}_{j=1}$
Rent Voucher	4.35	0.03	-0.02	-0.05	0.00	0.01	-0.00	-0.01	-0.00
Wage Subsidy	0.49	0.03	0.01	-0.05	0.00	-0.02	-0.01	-0.01	-0.00
Panel B. Parameters:									
	Baseline Welfare Gain	$\bar{A}_{j \neq 1}$	$\alpha_{I,j=1}$	ρ	$\alpha_{I,j \neq 1}$	λ	ω	Δ	θ_h
Rent Voucher	4.35	-0.03	-0.00	0.05	0.02	-0.12	-0.02	0.01	0.03
Wage Subsidy	0.49	-0.02	-0.00	0.05	-0.00	-0.04	0.02	-0.02	-0.00

Notes: This table provides an analysis of the sensitivity of welfare gains to changes in the parameter values used in our calibrated model. Columns 1-8 report results from increasing a given parameter (e.g., the altruism parameter $\tilde{\beta}$) by one percent. We examine sensitivity to changes in 16 different parameters spread across two panels of the table. Rows indicate whether the results are specific to either the wage-subsidy programs studied in Sections 5.1 and 5.2. For comparison, the left of the table reports the baseline welfare gains of 4.3 and 0.5 for the rent voucher and wage subsidy programs, respectively. See text for further details on all calculations.

B Estimation Details

B1 Child Skill Production Function

We rely on estimates from [Cunha et al. \(2010\)](#) for the calibrated model. Specifically, they estimate the following multistage production function for children's cognitive (c) and non-cognitive skills (nc):

$$\theta'_{q,k} = \left[\alpha_{1,q,j} \theta_{c,k}^{\rho_{q,j}} + \alpha_{2,q,j} \theta_{nc,k}^{\rho_{q,j}} + \alpha_{3,q,j} \theta_c^{\rho_{q,j}} + \alpha_{4,q,j} \theta_{nc}^{\rho_{q,j}} + \alpha_{5,q,j} I^{\rho_{q,j}} \right]^{1/\rho_{q,j}} e^{v_q}, \quad v_q \sim N(0, \sigma_{q,j,v})$$

for $q \in \{c, nc\}$. Using a nonlinear factor model with endogenous inputs, their main estimates, which are based on two-year periods, are reported in [Table B4](#) below. We interpret their first stage estimates as referring to the period in which the child is born in our model when the parent's age-period is $j = 8$ and the child's age-period is $j' = 1$, (i.e., 0–3 years old). The second stage is assumed to refer to the last period of skill development when the parent's age-period is $j = 11$ and the child's age-period is $j' = 4$ (i.e., 12–15 years old). We use linear interpolation to obtain the estimates for $j = 9$ and $j = 10$.

Table B4: Child Skill Production Function Estimates from [Cunha et al. \(2010\)](#)

	Cognitive Skills		Non-Cognitive Skills	
	1st Stage ($j = 8$)	2nd Stage ($j = 11$)	1st Stage ($j = 8$)	2nd Stage ($j = 11$)
Current Cognitive Skills ($\hat{\alpha}_{1,q,j}$)	0.479 (0.026)	0.831 (0.011)	0.000 (0.026)	0.000 (0.010)
Current Non-Cognitive Skills ($\hat{\alpha}_{2,q,j}$)	0.070 (0.024)	0.001 (0.005)	0.585 (0.032)	0.816 (0.013)
Parent's Cognitive Skills ($\hat{\alpha}_{3,q,j}$)	0.031 (0.013)	0.073 (0.008)	0.017 (0.013)	0.000 (0.008)
Parent's Non-Cognitive Skills ($\hat{\alpha}_{4,q,j}$)	0.258 (0.029)	0.051 (0.014)	0.333 (0.034)	0.133 (0.017)
Investments ($\hat{\alpha}_{5,q,j}$)	0.161 (0.015)	0.044 (0.006)	0.065 (0.021)	0.051 (0.006)
Complementarity parameter ($\hat{\rho}_{q,j}$)	0.313 (0.134)	-1.243 (0.125)	-0.610 (0.215)	-0.551 (0.169)
Variance of Shocks ($\hat{\sigma}_{q,j,v}$)	0.176 (0.007)	0.087 (0.003)	0.222 (0.013)	0.101 (0.004)

Notes: Standard errors in parentheses. The first stage refers to the period in which the child is born when the parent's age-period is $j = 8$ and the child's age-period is $j' = 1$ (i.e., 0–3 years old). The second stage refers to the period after the child is born when the parent's age-period is $j = 11$ and the child's age-period is $j' = 4$ (i.e., 12–15 years old).

To go from two-year periods to four-year periods (as in our model), we follow the steps in

Daruich (2020). Using $\hat{\alpha}$ to notate the estimates in Cunha et al. (2010) and α for the values in our model, the two main steps/assumptions for the transformation are: (i) we iterate in the production function under the assumption that the shock ν only takes place in the last iteration, i.e., replace $\theta_{q,k}$ by $\left[\alpha_{1,q,j} \theta_{c,k}^{\rho_{q,j}} + \alpha_{2,q,j} \theta_{nc,k}^{\rho_{q,j}} + \alpha_{3,q,j} \theta_c^{\rho_{q,j}} + \alpha_{4,q,j} \theta_{nc}^{\rho_{q,j}} + \alpha_{5,q,j} I^{\rho_{q,j}} \right]^{1/\rho_{q,j}}$;⁴⁷ and (ii) we assume that the cross-effect of skills (i.e., of cognitive on non-cognitive and of non-cognitive on cognitive) is only updated every two periods.⁴⁸ Under these assumptions, the persistence parameter needs to be squared (i.e., $\alpha_{1,c,j} = \hat{\alpha}_{1,c,j}^2$ and $\alpha_{2,nc,j} = \hat{\alpha}_{2,nc,j}^2$), while other parameters inside the CES function need to be multiplied by 1 plus the persistence parameter (e.g., $\alpha_{2,c,j} = (1 + \hat{\alpha}_{1,c,j}) \hat{\alpha}_{2,c,j}$).

B2 Replacement benefits: US Social Security System

The pension replacement rate is obtained from the Old Age Insurance of the US Social Security System. We use the skill level to estimate a proxy for average lifetime income, on which the replacement benefit is based. Average income at age j is estimated as $\hat{y}_j(\theta_c) = wE_j(\theta_c, \bar{\eta}) \times \bar{h}$ where $\bar{\eta}$ is the average shock (i.e., zero) and \bar{h} are the average hours worked (in the economy). Averaging over j allows average lifetime income $\hat{y}(\theta_c)$ to be calculated and used in (B1) to obtain the replacement benefits.

The pension formula is given by

$$\pi(\theta_c) = \begin{cases} 0.9\hat{y}(\theta_c) & \text{if } \hat{y}(\theta_c, e) \leq 0.3\bar{y} \\ 0.9(0.3\bar{y}) + 0.32(\hat{y}(\theta_c) - 0.3\bar{y}) & \text{if } 0.3\bar{y} \leq \hat{y}(\theta_c) \leq 2\bar{y} \\ 0.9(0.3\bar{y}) + 0.32(2 - 0.3)\bar{y} + 0.15(\hat{y}(\theta_c) - 2\bar{y}) & \text{if } 2\bar{y} \leq \hat{y}(\theta_c) \leq 4.1\bar{y} \\ 0.9(0.3\bar{y}) + 0.32(2 - 0.3)\bar{y} + 0.15(4.1 - 2)\bar{y} & \text{if } 4.1\bar{y} \leq \hat{y}(\theta_c) \end{cases} \quad (\text{B1})$$

where \bar{y} is approximately \$288,000 (\$72,000 annually).

⁴⁷We assume that the variance of the shock in the 4-year model is twice the one in the 2-year model (i.e., $\sigma_{q,j}, v^2 = \hat{\sigma}_{q,j}, v^2$).

⁴⁸Removing this assumption does not change results significantly since the weights corresponding to these elements are very small or even zero in the estimation (in Table B4, see row 2 under columns 1 and 2, as well as row 1 under columns 3 and 4), but it eliminates the CES functional form if $\rho_{c,j} \neq \rho_{nc,j}$.

C Welfare Measure

Our analysis centers on evaluating aggregate welfare under scenarios that feature different policies. Welfare is defined by the consumption equivalence under the veil of ignorance in the baseline economy relative to the economy with the policy in place. Formally, let $P \in \{0, 1, 2, \dots\}$ denote the set of policies, with $P = 0$ being the initial economy (with no voucher or wage-subsidy program) in steady state. We refer to the consumption equivalence as the percentage change in consumption Δ in the initial economy that makes individuals indifferent between being born in the initial economy ($P = 0$) and the one in which the policy $P \neq 0$ is in place. Denote $V_{j=5}^P(a, \theta, n, \epsilon, \Delta)$ be the welfare of agents with initial state of the economy if their consumption (and that of their descendants) were multiplied by $(1 + \Delta)$:

$$\widetilde{V}_{j=5}^P(a, \theta, n, \epsilon, \Delta) = \mathbb{E}^P \sum_{j=5}^{j=J_d} \beta^{j-5} u(c_j^P(1 + \Delta), h_j^P, n_j) + \beta^{12-5} \delta \widetilde{V}_{j'=5}^P(\widehat{a}, \theta_k, n_{j=11}, \epsilon', \Delta),$$

where for the sake of clarity the expression above suppresses the utility terms for moving costs and disutility of time with children. Note that the policy functions are assumed to be unchanged when Δ is introduced. For example, consumption c_j^P is consumption chosen by individuals in economy P (in age j) and is not affected by Δ . For any measure Δ , the average welfare is:

$$\overline{V}^P = \int_{a, \theta, n, \epsilon} \widetilde{V}^P(a, \theta, n, \epsilon, \Delta) \mu_p(a, \theta, n, \epsilon),$$

where μ_p is the distribution of initial states $\{a, \theta, n, \epsilon\}$ in the economy P . We define Δ^P as the consumption equivalence that makes individuals indifferent between being born in the baseline economy or one in which policy P is in place:

$$\overline{V}^0(\Delta^P) = \overline{V}^P(0).$$

By definition, the welfare gains come from two sources. First, there are changes in the expected discount utilities at each state $\widetilde{V}_{j=5}^P(a, \theta, \epsilon, n, \Delta)$. Second, there are also shifts in the probabilities of each state $\mu_p(a, \theta, n, \epsilon)$.