Public Education and Intergenerational Housing Wealth Effects

Michael Gilraine, James Graham, and Angela Zheng*

December 2, 2025

Abstract

While rising house prices are known to benefit existing homeowners, we document a novel channel through which house price shocks have intergenerational wealth effects. Using panel data from school zones within a large U.S. school district, we find that higher local house prices lead to improvements in local school quality, thereby increasing children's human capital and future incomes. We then study housing market shocks in an overlapping generations model with neighborhood choice, spatial equilibrium, and endogenous school quality. We find that the school quality channel generates large intergenerational wealth effects that account for around one-third of total housing wealth effects.

Keywords: House prices, intergenerational wealth effects, school quality, neighborhood choice, intergenerational mobility

JEL: E21, E24, I24, J62, R21, R23.

*Michael Gilraine: Simon Fraser University, gilraine@sfu.ca. James Graham: University of Sydney, james.a.graham@sydney.edu.au. Angela Zheng: McMaster University, zhenga17@mcmaster.ca. We thank Victoria Gregory, Matthijs Korevaar, Julian Kozlowski, Oksana Leukhina, Suresh Naidu, B. Ravikumar, Hannah Rubinton, Thomas Sargent, Diego Daruich, Luiz Brotherhood, Kyle Herkenhoff, and seminar participants at the NBER Summer Institute, the Society for Economic Dynamics meetings, the BSE Summer Forum, the Opportunity & Inclusive Growth Institute, the Mid-West Macroeconomics Conference, the North American Meeting of the Urban Economics Association, the Society for Nonlinear Dynamics Symposium for Young Researchers, the Canadian Labour Economic Forum, the Virtual Australian Macroeconomics Seminar Workshop, the New Zealand Association of Economists meetings, the Federal Reserve of St. Louis, the University of Houston, Texas A&M University, Tokyo University, Hitotsubashi University, Melbourne University, Australian National Unviersity, and the University of Queensland. Angela Zheng thanks the Social Sciences and Humanities Research Council of Canada for funding this project through an Insight Development Grant. Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at `http://www.zillow.com/ztrax`. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

1. Introduction

Recent research has documented strong links between intergenerational income mobility and the neighborhood choices of parents.[1] These studies show that where a family chooses to live shapes children's future earnings through local factors such as resident composition and school quality.[2] But access to good neighborhoods can be expensive as the value of these local factors is capitalized into house prices.[3] As a result household wealth, through housing, also determines spatial patterns of inequality.

In this paper we identify a new mechanism through which housing market shocks can reshape and propagate economic opportunity across generations. We show that rising local house prices improve the quality of local public schools, which leads to higher children's future earnings. We characterize this mechanism as an intergenerational housing wealth effect.

We proceed in two stages. First, we provide empirical evidence that local house price growth leads to improvements in local public school quality as measured by school value-added. Second, we build an overlapping generations model that incorporates neighborhood choices, local schools, and spatial equilibrium. Our model enables us to analyze the intra- and intergenerational transmission of permanent housing market shocks. We demonstrate that intergenerational wealth effects due to changing school quality account for one-third of total housing wealth effects. While our empirical findings establish an endogenous relationship between house prices and school quality, the model allows us to investigate its broader implications in the absence of detailed intergenerational data.

Empirically, we show that local house price growth leads to changes in household sorting across neighborhoods, increasing the share of high socioeconomic status residents. This demographic change attracts higher-quality teachers which improves local school quality. Importantly, these improvements are independent of changes in school financing mechanisms such as local property taxes.[4] Throughout both our empirical analysis and modeling work, we argue that higher school quality contributes to greater human capital formation among resident children, thereby increasing their future earnings. Consequently, families exposed to positive local housing market shocks benefit from both contemporaneous increases in housing wealth and from higher children's incomes. These outcomes represent both intra- and intergenerational wealth effects of housing market shocks.

Our first contribution is an empirical analysis combining student-level administrative

[1]See, for example, Chetty et al. (2014c) and Chetty et al. (2016).

[2]Recent modeling work in this area includes work by Fogli et al. (2025), Zheng et al. (2021), Gregory et al. (2022), and Chyn et al. (2025).

[3]See, for example, Black (1999).

[4]For papers studying the link between property taxes and inequality, see Benabou (1994), Benabou (1996a), Benabou (1996b), Durlauf (1996a), Durlauf (1996b), Fernández et al. (1996), Fernández et al. (1998), and Davis et al. (2022).

data from a large, urban, U.S. school district and local house prices constructed from housing transactions data. We document that faster house price growth in a school attendance zone leads to larger subsequent improvements in local public school quality, as measured by school value-added. To alleviate concerns about endogeneity and measurement error, we present estimates using two instruments for house prices: (i) a shift-share instrument that exploits geographic variation in the composition of local housing characteristics, and (ii) local housing supply elasticities.[5] We find that a 100 percent increase in house prices over a five-year period leads to a 0.25-standard deviation increase in local school value-added. Drawing on work by Chetty et al. (2014b), we provide a back-of-the-envelope calculation linking these changes in school quality to lifetime gains in child incomes. This exercise suggests that a one standard deviation increase in house price growth is associated with a \$37,000 increase in the present-value of lifetime income in year 2000 dollars.

We also explore mechanisms that might account for the empirical relationship between house price growth and school quality changes. First, note that the relationship cannot be explained by changes in local school financing. This is because we study school zones within a school district, while residential property taxes accrue to the broader school district. Second, we show that faster house price growth alters the composition of local schools by reducing the share of low socioeconomic status students and increasing the share of students with college-educated parents. Third, we show that improving teacher quality due to teachers switching schools accounts for around two-thirds of the overall change in school quality. We argue that this sorting of teachers across schools is likely to be linked to the observed changes in student body composition.[6]

Our second contribution is to develop an overlapping generations model that quantifies intra- and intergenerational housing wealth effects in the presence of a semi-endogenous school quality channel. In the model, young households first choose a neighborhood in which to live and purchase a house in. Once children arrive, parents have the opportunity to move by selling their existing house and buying a new one in a different neighborhood. Neighborhood choice is of central importance because it determines the local schools that children attend, and school quality is a significant input into human capital formation. Parents value their children's lifetime wealth, which comprises direct transfers of wealth and the present discounted value of their future income.

In the spatial equilibrium of the model, house prices adjust to clear neighborhood housing markets, and school qualities adjust to reflect relative house prices. A crucial parameter in the model is the elasticity of school quality to house prices, which we take directly from our empirical estimates. The remaining model parameters are calibrated to capture cross-neighborhood differences in: house prices, household incomes, average

[5] See Graham et al. (2023) and Baum-Snow et al. (2024), respectively.

[6] See, for example, Bonhomme et al. (2016).

student test scores, and intergenerational income mobility (see Chetty et al., 2018).

We then study intra- and intergenerational housing wealth effects in the model via exogenous, permanent neighborhood demand shocks.[7] Along an equilibrium transition path to the new steady state, positive local demand shocks increase house prices, which increase local school quality in turn. Households that purchased homes in these neighborhoods prior to the shock may then adjust consumption, transfers to children, or investments in school quality through decisions to stay in or move across neighborhoods.

Our primary focus is on the marginal change in children's lifetime incomes with respect to permanent house price shocks, which we estimate to be 1.7 cents in the dollar. Combining this estimate with a \$130,000 rise in house prices – equivalent to the observed standard deviation of 5-year house price changes – we compute an average increase in children's life-time incomes of \$34,000. This estimate is very similar to the \$37,000 income gain we infer from our empirical results. We also report an annualized marginal propensity to consume out of house prices of 2.4 cents in the dollar, which falls within the range of recent estimates in the literature (Mian et al., 2013; Aladangady, 2017; Graham et al., 2023). Additionally, the average annualized marginal propensity to transfer wealth to children is 1.2 cents in the dollar, consistent with recent empirical work on intergenerational transfers of housing wealth shocks (Daysal et al., 2023). Taken together, intergenerational effects due to higher children's incomes and direct transfers account for over half of total housing wealth effects.

Finally, we show that intergenerational housing wealth effects are the result of active parental choices to invest housing gains in additional education. Households benefiting from positive housing market shocks are somewhat more likely to stay in their current neighborhoods or move to more expensive ones with better schools. We find this mechanism is particularly important for low-income families. In contrast, households facing negative housing market shocks are more likely to leave their current neighborhoods but tend to move to areas with lower house prices and lower school quality. Our results suggest that parents view education quality as an investment in their children's human capital and leverage housing wealth gains to enhance these investments. This education investment mechanism is reinforced by our endogenous school quality channel. These strong connections between housing wealth, school quality, and neighborhood choice are central to the intergenerational transmission of housing wealth shocks.

1.1. Related Literature

This paper follows a literature studying the relationship between intergenerational inequality, neighborhood choice, school quality, and child human capital accumulation

[7] Note that since we only model homeowners, our housing wealth effects are conditional on homeownership. The homeownership rate in the state containing our school district of interest was 58 percent in the mid-2000s (U.S. Census Bureau, 2019).

(Benabou, 1994; Benabou, 1996a; Benabou, 1996b; Durlauf, 1996a; Durlauf, 1996b; Fernández et al., 1996; Fernández et al., 1998). Many of these papers focus on the link between local property taxes and school financing across school districts. In this context, house price changes directly impact school revenues and thus quality, as documented by Davis et al. (2022) and modeled in Zheng et al. (2021). In contrast, we study school zones within a school district, whereas property taxes are collected at the district level.[8] This allows us to isolate differences in school quality due to local factors such as the composition of students or the quality of teachers at these schools.

Our model builds on recent work studying intergenerational inequality with neighborhood sorting and endogenous local school quality (Kotera et al., 2017; Aliprantis et al., 2018; Fogli et al., 2025; Eckert et al., 2019; Zheng et al., 2021; Gregory et al., 2022; Chyn et al., 2025). The most closely related research to our own is Zheng et al. (2021), Fogli et al. (2025), and Chyn et al. (2025). These papers build similar overlapping generations models with neighborhood choice, endogenous sorting, and local spillovers into child human capital accumulation. Fogli et al. (2025) study a permanent increase in the skill premium that encourages additional human capital investment. Their shock increases neighborhood segregation along income lines and helps explain increasing dispersion of cross-neighborhood intergenerational income mobility since the 1980s. Both Zheng et al. (2021) and Chyn et al. (2025) study dynamic equilibrium responses to policy changes such as the introduction of school vouchers, transfers, or place based-subsidies. In contrast, our paper studies wealth effects following neighborhood-specific housing demand shocks. Our results highlight that even generic fluctuations in house prices can have large effects on intergenerational mobility and the transmission of wealth. Another key contribution of our work is that we explicitly model homeownership. In contrast, prior literature largely models renters with absentee landlords. The benefit of our approach is the ability to study interactions between housing wealth and inequality, as well the wealth effects of housing market shocks.

Our research is also related to the empirical literature estimating contemporaneous wealth effects of house price changes on current homeowners (Mian et al., 2013; Aladangady, 2017; Graham et al., 2023). In recent work, Daysal et al. (2023) use Danish administration data to study the intergenerational transmission of wealth via house price shocks to parent homeowners. Benetton et al. (2022) use U.S. credit records to show that home-owning parents respond to housing wealth shocks by extracting home equity to provide children with the resources to access their own first homes. Relatedly, Brandsaas (2025) builds a life-cycle housing model to study how transfers of wealth to adult

[8] In any case, property tax revenues in our school district of interest account for just fifteen percent of total school revenues. The only other source of funding that may vary with local house prices are donations from parent-teacher associations and school booster clubs. However, these account for just 0.4 percent of funding in our district, and around one percent of aggregate spending on education in the U.S. (Brown et al., 2017).

children help them enter the housing market. In our paper we jointly study intra- and intergernational housing wealth effects in the context of local housing market shocks with a semi-endogenous local school quality response.

Finally, our paper connects to the education economics literature linking school quality to student body composition (Rothstein, 2006; Allende, 2019). We provide new evidence that teacher sorting is a key driver of the relationship between school quality and student demographics. This sorting is consistent with the view that teachers prefer higher achieving students, which is supported by evidence on teacher preferences across school assignments (Boyd et al., 2011; Bonhomme et al., 2016; Johnston, 2020; Karbownik, 2020).

2. Empirical Analysis

2.1. Data

Education Data: We use administrative data from a large urban school district in the United States. The data cover all students and teachers in public schools in the district for academic years 2003-04 through 2016-17. We observe mathematics and English test scores on standardized end-of-grade exams for each student in each year of schooling, with the exception of 2013-14 when no testing was conducted. The data also provide demographic information for each student. Since our interest is in the relationship between residential location and school quality, and because out-of-zone school choices are much more readily available for high school students, we restrict our sample to students in grades K-5.

Regarding external validity, our school district is broadly similar to others in the U.S. Teachers in the district are paid according to fixed salary schedules, as in 89 percent of school districts in the country (Hansen et al., 2017). Annual teacher turnover rates in the district are comparable to the nationwide average of 16% (Carver-Thomas et al., 2017). And while we focus on public schools, we note that our district's private school share is 8%, similar to the nationwide average of 10% (Snyder et al., 2012).

For each elementary school in the district, we construct a measure of school quality called value-added (VA) using standard methods in the economics of education literature.[9] To do this, we prepare the data by first normalizing student test scores within each grade and year to have zero mean and unit variance. Since we require both current and lagged test scores to construct VA, we exclude all students with invalid scores in the current or previous year, and we exclude data from 2013-14 and 2014-15 due to the lack of testing in 2013-14. Our final sample consists of 1.6 million student-year observations covering

[9] School VA methods rely on the assumption that student assignment to schools is uncorrelated with unobserved determinants of achievement, conditional on controls. Crucially, these controls include lagged student test scores. See Deming (2014) for validation of these measures.

around 700,000 unique students across 420 elementary schools. Appendix A.1 provides additional details and summary statistics.

To estimate value-added, we first regress student test scores on school fixed effects and observable determinants of student performance. These controls are: (i) year and grade dummies, (ii) cubic polynomials in students' prior-year test scores in mathematics and English, each interacted with grade dummies, and (iii) student-level demographics, including parental education, economically disadvantaged status, ethnicity, gender, limited English status, and age, all interacted with grade dummies. School-year fixed effect estimates are then given by the average of students' residualized test scores at a given school in a given year. We then shrink the estimated fixed effects using an empirical Bayes method (see Morris, 1983) since the raw fixed effect estimates overstate the variance of school VA (Koedel et al., 2015). Appendix A.2 describes our VA estimation procedure in more detail.

Our procedure produces VA estimates for each school-year combination. To interpret the VA measure, note that students moving to a school with a one-unit increase in VA would be expected to score one-standard deviation higher in the overall student test score distribution.

House Price Data: The ZTRAX database provides transaction-level housing data for the US state that contains our school district of interest (Zillow, 2020). We use these data to construct real annual house price indexes for each school zone within our school district for academic years 1998-99 to 2018-19.

The address of each house sold within our district is matched to a school zone using the latitude-longitude coordinates of the property. Since school zone boundaries may change over time, we use school zone shapefiles from 2008-09 (The College of William and Mary and the Minnesota Population Center, 2011) and 2015-16 (National Center for Education Statistics, 2018). Approximately 8 percent of houses cannot be matched to a school zone or change zones across years, and we exclude these houses from our sample. Our final sample covers 393 school zones with at least thirty house sales per year.

We deflate all nominal house values by the CPI (U.S. Bureau of Labor Statistics, 2021) and then construct an arithmetic repeat-sales house price index following Shiller (1991). In contrast with a median sales price index, the repeat-sales index holds constant the composition of the housing stock over time. Table B.2 in Appendix C presents summary statistics for our housing data.

School Zone Demographic Data: We gather information on school zone-level sociodemographic characteristics from the American Community Survey (ACS) 5-year estimates (U.S. Census Bureau, 2019). This demographic information includes average educational attainment, homeownership rates, and family structure. Since ACS data are not available for school zones, we construct a cross-walk between census tracts and school zones. The cross-walk aggregates census tract-level demographics to the school zone level using census tract-level population weights. Appendix A.5 provides additional details. Table B.2 in Appendix C reports summary statistics on sociodemographic characteristics for the average school zone in our sample.

2.2. Empirical Strategy

We estimate the relationship between changes in house prices and subsequent changes in school quality using the following regression specification:

$$\Delta \text{VA}_{z,t,t+5} = \alpha_z + \alpha_t + \beta \Delta \log HousePrices_{z,t-5,t} + \delta' X_{z,t,t+5} + \epsilon_{z,t} \tag{1}$$

where $\Delta \text{VA}_{z,t,t+5}$ is the change in school VA in school zone z between years t and $t+5$, and $\Delta \log HousePrices_{z,t-5,t}$ is the lagged change in the log of the repeat-sales index in school zone z between years $t-5$ and t. Our coefficient of interest is β, which measures the elasticity of school VA with respect to local house prices. The vector $X_{z,t,t+5}$ includes controls for sociodemographic characteristics in school zone z measured between the years t and t+5, such as the homeownership rate and the share of married families with children. School zone fixed effects, α_z, account for school-specific factors affecting average school quality growth. For example, schools with good reputations may improve over time at a faster rate than others. Time fixed effects, α_t, absorb common trends across school zones such as broader economic forces affecting the entire school district. Thus, our regression specification exploits relative house price changes across school zones within the school district. Throughout our empirical analysis, we cluster standard errors at the school zone level.

Our baseline regression in Equation (1) makes two assumptions about the dynamics of the relationship between school quality and house prices. First, changes in house prices affect school quality with a lag. Second, these changes take place over several years. Both assumptions reflect our view that it takes time for changes in house prices to affect local schools. In Section 2.4, we report the sensitivity of our results to alternative timing assumptions.

To estimate Equation (1) we use house price data from 1998-99 to 2018-19, while our VA measure is available from 2003-04 to 2016-17 excluding the years 2013-14 and 2014-15. After constructing 5-year growth rates, our sample consists of years $t \in$ {2003-04, 2004-05, 2005-06, 2006-07, 2007-08, 2010-11, 2011-12 }. As motivating evidence, Figure 14 in Section C presents a binscatter plot of the relationship between $\Delta \log HousePrices_{z,t-5,t}$ (in percentiles) and $\Delta VA_{z,t,t+5}$, residualized against year and school zone fixed effects. We now turn to discussing causal estimation of Equation (1).

2.3. Identification

We face two challenges to identification in estimating Equation (1). First, house price growth may be correlated with unobserved variables in the error term $\epsilon_{z,t}$ that are themselves correlated with changes in local school quality. For example, improvements in local amenities could induce higher demand for housing at the same time as predicting higher future school quality. Second, there may be measurement error in house price growth since we only observe the sample of houses that happen to sell in each school zone in a given year.

To address these concerns, we use an instrumental variable estimation strategy employing two instruments from the recent housing literature: (i) a Bartik-style or shift-share instrument following Graham et al. (2023), and (ii) local housing supply elasticities from Baum-Snow et al. (2024). The first instrument exploits geographic variation in the composition of the housing stock given aggregate changes in the demand for different kinds of housing. The second instrument exploits geographic variation in the ease of constructing new housing given changes in aggregate demand for all types of housing.

Bartik-Style House Price Instrument: Following Graham et al. (2023), the Bartik-style instrument is constructed by taking the local share of houses with given physical characteristics and interacting those shares with aggregate estimates of the marginal prices of those characteristics. For example, we combine the share of two-bedroom houses in each school zone with the aggregate marginal price of two-bedroom houses. Again making use of the ZTRAX data, we proceed in two stages. First, we use three house characteristics that are widely reported in the data: decade of construction, number of bedrooms, and number of bathrooms. We compute the local shares of houses possessing each characteristic using the set of unique properties sold between 1999 and 2019.

Second, we estimate the aggregate marginal prices of each characteristic with a hedonic house price regression. The regression features time-varying coefficients on a set of dummy variables capturing our chosen characteristics. For example, bedroom characteristics are represented by dummy variables for houses with 1, 2, 3, 4, or 5+ bedrooms. The growth rate in the marginal price of a given house characteristic is the change in the estimated time-varying coefficient on the associated dummy variable. To capture aggregate marginal prices, we use transactions for all houses in the state containing our school district, but exclude transactions from the school district itself. This is similar to the leave-one-out estimator used for shift-share instruments, expcept that we exclude all sources of variation in house prices that might directly affect school zones in our district (i.e., all other zones within the district). For further details on instrument construction, see Appendix A.6.

Identification of Equation (1) using our Bartik-like instrument follows from exogeneity of the local housing characteristics shares (Goldsmith-Pinkham et al., 2020). Specifically, cross-sectional variation in local house characteristics must be exogenous to the error term $\epsilon_{z,t}$. In other words, unobserved shocks to local school quality must be uncorrelated with the composition of the local housing stock. We think this is plausible because house characteristics are largely predetermined at the time of other shocks affecting local school quality. Table B.3 in Appendix C shows that the composition of the housing stock is extremely persistent. The transaction-weighted average of 15-year within-zone correlations for our characteristics shares is 0.84 for number of bedrooms, 0.88 for number of bathrooms, and 0.94 for decade of construction. Since it takes time for new residential construction to affect the composition of the total housing stock, it seems likely that housing characteristics are unresponsive to short- or medium-run shocks affecting the quality of local schools.

Housing Supply Elasticity Instrument: Our second instrument uses local housing supply elasticities from Baum-Snow et al. (2024).[10] Baum-Snow et al. (2024) identify these elasticities by combining an urban economic geography model with estimated changes in housing demand due to labor income shocks across nearby commuting destinations. As suggested by the authors we use estimates from their quadratic finite mixture model, we use their supply elasticities for housing units, and we use their estimates for 2001 which occurs prior to the start of our own sample period. Since they provide estimates for census tracts, we again make use of our cross-walk to aggregate up to the school zone level (see Appendix A.5). In order to produce time series variation in the instrument, we interact the cross-sectional housing supply elasticities with the aggregate 5-year growth rate of real house prices for the state in which our school district is located.[11]

Identification of Equation (1) using the housing supply elasticity instrument requires that unobserved shocks to local school quality are uncorrelated with the sensitivity of local housing provision to house price changes. As discussed in Gyourko et al. (2008), Saiz (2010), and Baum-Snow et al. (2024), elasticities of housing supply are functions of local land topography and local land use regulations. Existing features of the local landscape, such as the presence of local water features or whether the land is on an incline, are almost certainly unrelated to the growth rate of local school quality. Local regulations could be related to changes in school quality to the extent that local politics influence both regulation and school policies. However, both land use and school policies are generally determined at higher levels of geography, such as the city or school district, rather than at the level of the local school zone.

[10]This follows seminal work by Saiz (2010) who produces housing supply elasticities for larger Metropolitan Statistical Areas.

[11]Aladangady (2017) interacts housing supply elasticities with changes in national interest rates, and Graham et al. (2023) interacts housing supply elasticities with broad, regional house price changes.

2.4. The Effect of House Prices on School Quality

Table 1 presents our estimates of the effect of house prices on school quality from Equation (1). Columns (1)–(2) report OLS estimates, Columns (3)–(4) report 2SLS estimates using the Bartik-style instrument, Columns (5)–(6) report 2SLS estimates using the housing supply elasticity instrument, and Columns (7)–(8) report 2SLS estimates from an overidentified specification using both instruments. Each specification is estimated first with year fixed effects only, and then with a full set of fixed effects and controls.

Table 1: Effect of House Price Growth on School Value Added

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ House Price	0.088***	0.126***	0.212***	0.250***	0.291	0.351*	0.217***	0.253***
	(0.028)	(0.030)	(0.062)	(0.063)	(0.187)	(0.211)	(0.060)	(0.062)
School Zones	393	393	393	393	393	393	393	393
Specification	OLS	OLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Instrument	–	–	G&M	G&M	BS&H	BS&H	Both	Both
1st Stage F-Stat	–	–	174.10	212.23	11.17	9.25	94.81	112.5
Sargan Stat.	–	–	–	–	–	–	0.45	0.75
Zone F.E.	No	Yes	No	Yes	No	Yes	No	Yes
Year F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Zone Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1,873	1,873	1,873	1,873	1,873	1,873	1,873	1,873

Notes: This table presents estimates of Equation (1). Columns (1) and (2) are estimated via OLS, Columns (3) and (4) are IV estimates using the Shift-Share instruments, Columns (5) and (6) are estimates using the BSH instrument, and Columns (7) and (8) use both instruments. School zone controls include: homeownership rate, percentage of individuals with a bachelor's degree or higher, and share of families that are married with children. Standard errors are clustered at the school zone-level and are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01

Our OLS specification produces statistically significant estimates in the range of 0.088 to 0.126. Using our Bartik-style instrument for house prices, we estimate larger coefficients of 0.212 to 0.250. First-stage F-statistics for these regressions are in the range of 174 to 212, indicating a strong relationship between our instrument and house price growth. Using our housing supply elasticity instrument, we estimate slightly larger coefficients of 0.291 to 0.351. These estimates are statistically noisy, and only the specification with all controls is statistically significantly different from zero at the 10 percent level. Additionally, the housing supply elasticity instrument is much weaker than the Bartik-style instrument: first-stage F-statistics are in the range of 9 to 11.[12] Finally, our overidentified specification with both instruments produces statistically significant estimates of 0.217 to 0.253, which is very similar to our estimates with the Bartik-style instrument alone. Under this specification we conduct overidentification tests, which produce Sargan statistics of 0.45 to 0.75. Given a critical value at the 5% level of 3.84, we cannot reject the null hypothesis that the instruments are jointly valid instruments.

Our preferred estimate in Column (8) indicates that 100 percentage point faster house

[12]This is consistent with Graham et al. (2023) who also find that housing supply elasticity instruments are weak predictors of house price growth in a panel data context.

price growth rate is associated with a 0.253 standard deviation increase in school VA. This is the same as 25 percent of a standard deviation gain in average student test scores.

To provide an economic interpretation of our estimates, we conduct a back-of-the-envelope calculation to translate the increase in school VA and student test scores into future income gains. First, a one standard deviation increase in house price growth (65 percentage points) is associated with 0.16 ($= 0.253 \times 0.65$) of a standard deviation increase in student test scores in each year of schooling. Second, Chetty et al. (2014b) report that a standard deviation increase in test scores during a single grade year is associated with a present value gain in lifetime income of $38,950 in 2000 dollars. Therefore, the initial house price shock is associated with a lifetime income gain of $6,232 ($= \$38,950 \times 0.16$) for each year of schooling. Finally, a child that completes six years of elementary schooling can expect lifetime income gains of $37,392 (=$6,232×6) following a standard deviation shock to house prices in their school zone.

Finally, we illustrate the sensitivity of our results to our choice of 5-year growth rates in house prices and school qualities for estimating Equation (1). Table B.4 in Appendix C estimates the effect of house price changes on school VA over 3-, 4-, 5-, 6-, and 7-year horizons. Our estimates are always statistically significant and monotonically increasing with the length of adjustment horizon. Our 3-year estimates are as small as 0.14, while our 7-year estimates are as large as 0.389. These results emphasize that any effect of house price changes on school quality is likely to take place over the medium- to long-run.

2.5. Mechanisms

We now investigate the mechanisms by which house price growth could lead to improvements in school quality. Since we study changes across school zones within a school district, differences in property tax revenues cannot explain this relationship. Additionally, we do not think that changes in local school or parent resources associated with house price growth are directly associated with higher quality schooling. Column (4) of Table 2 shows that house price shocks do not lead to changes in average class sizes, suggesting that local schools do not receive additional funding that could be used to hire more teachers. Column (5) shows that house price shocks do not lead to changes in private school enrollment shares, suggesting that house price rises do not induce additional parent spending on private education.

Instead, we consider the way in which house price growth affects the composition of students in the local school zone. Column (1) of Table 2 shows that faster house price growth leads to a sizeable reduction in the share of free and reduced-price lunch students. Column (2) shows that the share of students whose parents have college education increases. And Column (3) indicates that higher house prices may also reduce the share of visible minority students, though the estimate is statistically insignificant.

Table 2: Effect of House Prices on School Characteristics

	Share FRL	Share College	Share Black	Class Size	Private/Public
	(1)	(2)	(3)	(4)	(5)
Δ log House Price	-0.569^{***}	0.053^{*}	-0.013	0.100	0.008
	(0.091)	(0.028)	(0.009)	(0.928)	(0.029)
School Zone F.E.	Yes	Yes	Yes	Yes	Yes
Year F.E.	Yes	Yes	Yes	Yes	Yes
First-Stage F-stat	96.5	107.93	93.4	89.7	98.1
Number of Schools	396	393	394	392	396
Observations	2,203	1,873	2,680	2,372	2,683

Notes: This table presents coefficients on Δ log House Price estimated from Equation (1) with different dependent variables. In Column (1) the dependent variable is ΔFRL, the change in the share of free and reduced lunch students from t to $t+5$. Column (2) looks at the change in students whose parents have college education. Column (3) estimates the effect of Δ log House Price on $\Delta Black$, the change in the share of Black students in a school zone from t to $t+5$. In Column (4) the dependent variable is the change in average class size in the school zone, while in Column (5) it is the ratio of students in private to public school. Private school ratio is defined as the number of students attending a private school in a zone over the number of students attending the public catchment school. All estimates are computed via 2SLS using the shift-share and BSH instruments. *p<0.1; **p<0.05; ***p<0.01

Peer- and Peer-Invariant Value Added: We now test whether faster house price growth is associated with improving school quality (i) directly through peer effects, or (ii) indirectly through changes in the quality of instruction (Rothstein, 2006; Allende, 2019). To explore these channels, we follow Altonji et al. (2015) and Allende (2019) by decomposing school VA into the contributions of the student body (i.e. peer VA) and the contributions of non-peer inputs into school quality such as teachers, principals, class size, infrastructure, and curriculum (i.e. peer-invariant VA).

We give a brief overview of the methodology here and relegate a detailed description to Appendix A.3. The peer component of school VA is derived from student characteristics that may affect the outcomes of other students. We follow Allende (2019) in assuming that these characteristics are well-represented by socioeconomic status and parents' education level. We then project school-year VA onto these peer characteristics plus a school fixed effect. Peer VA is then given by the relationship between year-to-year variation in school VA and year-to-year changes in peer characteristics. The residual component of school VA is labelled peer-invariant VA.

Table 3 presents 2SLS estimates of the effect of house prices on the two components of school VA. Column (1) repeats our preferred estimate of the effect of prices on school VA from Table 1. Column (2) reports the effect of changes in house prices on the peer component of VA, with an estimated coefficient of 0.010 that is statistically significant from zero. Column (3) reports the effect of changes in house prices on the peer-invariant component of value-added, with an estimated coefficient of 0.243. These results suggest that peer VA accounts for just 4 percent of the change in total VA following a house price shock, while the remainder is due to changes in the peer-invariant component of VA. Our finding of a small direct impact of school peers on school quality is consistent with a

Table 3: Effect of House Prices on Peer, Peer-invariant and Teacher VA

	Δ School VA	Δ PeerVA	Δ Fixed VA	Δ Teacher VA
	(1)	(2)	(3)	(4)
Δ House Price	0.253***	0.010***	0.243***	0.184***
	(0.062)	(0.0027)	(0.061)	(0.047)
School Zones	393	393	393	393
First-Stage F Stat	112.52	112.52	112.52	111.9
School Zone F.E.	Yes	Yes	Yes	Yes
Year F.E.	Yes	Yes	Yes	Yes
School Zone Controls	Yes	Yes	Yes	Yes
Observations	1,873	1,873	1,873	1,867

Notes:This table presents estimates of Equation (1), where the dependent variable is replaced with different measures of school value added. Column (1) estimates effects on total school VA. Column (2) estimates effects on the peer component of school VA. Column (3) estimates effects on the peer-invariant component of school VA. Column (4) estimates effects on teacher VA through changes in teacher quality induced by teacher entry and exit. All columns report 2SLS estimates using the shift-share and BSH instrument. Standard errors and first stage F-statistics are clustered at the school zone-level and standard errors are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01

large literature documenting fairly modest effects of peers on child outcomes (Sacerdote, 2011).

Teacher Quality: We now investigate the extent to which changes in peer-invariant VA is due to changes in teacher quality. We are interested in identifying changes in teacher quality due to teacher sorting across schools rather than due to within-teacher quality changes at a given school. Within-teacher quality changes may be difficult to disentangle from changing student demographics if teachers find it easier or more rewarding to teach advantaged students. We therefore follow the teacher-switching literature (Chetty et al., 2014a; Bacher-Hicks et al., 2014; Gilraine et al., 2021) and compute changes in teacher VA at each school due solely to changes in staff (i.e., teacher entry to and exit from a particular school). Note that teacher movements across schools are fairly common: Table B.5 in Appendix C reports one- and five-year teacher turnover rates in our school district of 20 and 50 percent, respectively.

We first estimate teacher-level VA for each teacher and year in our data set.[13] We then compute changes in school VA using a jackknife procedure that excludes data for teachers while they remain at a given school but includes data for teachers when they switch schools. Changes in a school's VA are then computed from the enrollment-weighted means of additions of teacher VA when new teachers enter and subtractions of teacher VA when old teachers exit. Appendix A.4 provides details of the estimation.

Column (4) of Table 3 reports our estimate of the effect of house prices on school quality through teacher switching. We estimate a statistically significant coefficient of 0.18. Comparing Columns (3) and (4), our results suggest that turnover-induced changes in teacher quality account for three-fourths of peer-invariant value-added. The remaining small changes in peer-invariant VA are due to other school-specific factors such as within-

[13]This is similar to our estimates of school VA, and follows standard procedures in the VA literature. Details are provided in Appendix A.4.

teacher improvements, better matching between students and teachers, higher quality principals, and changes in school curricula.

Our results suggest that house price growth largely drives school quality improvements through a teacher sorting mechanism. Note that there are rigid teacher salary schedules within the school district that preclude cross-zone wage adjustments that might attract better teachers (see Hansen et al., 2017) and changes to school zone house prices do not affect teachers' real wages since most teachers commute from outside their school zone (Arturo Santelli et al., 2022). Instead, previous work has shown that teachers sort across schools on the basis of student composition (Rothstein, 2015; Bonhomme et al., 2016; Karbownik, 2020; Bates et al., 2022). Since teachers prefer working with students from advantaged backgrounds (Allensworth et al., 2009; Boyd et al., 2011), they likely respond to house price shocks by moving to schools with an increasing share of high socioeconomic status students. This is consistent with the results in Table 2. Note that teachers will value this amenity even in the absence of peer effects, since it improves their workplace experience even if it does not influence workplace productivity.

3. Model

Our empirical results show that rising local house prices are associated with improvements in local school quality. We now build a model to study the intra- and intergenerational wealth effects of house price shocks in the presence of a semi-endogenous local school quality channel.

3.1. Environment

Overview: The model features overlapping generations of parent-child households. Households live for four periods, where age is denoted $j \in \{0, 1, 2, 3\}$. Figure 1 summarizes the timeline of life events including schooling, earnings, consumption, borrowing and saving, leaving transfers to children, and choosing neighborhoods in which to live. The desirability of each neighborhood varies with the price of housing, school quality, common neighborhood preference shifters, and idiosyncratic preference shocks. In equilibrium, house prices adjust to clear local housing markets and school quality is a function of local house prices.

Figure 1: Timeline of Household Events and Decisions

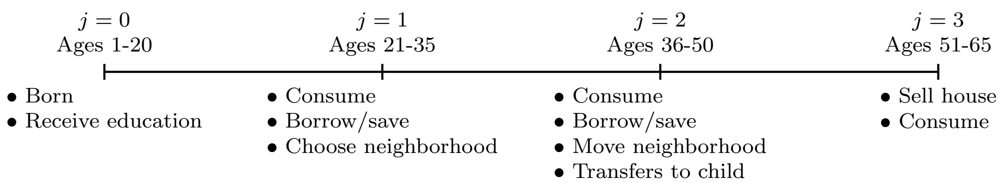

Neighborhoods and Housing: There are five model neighborhoods denoted $n \in \{A, B, C, D, E\}$. Households purchase one unit of housing at a neighborhood-specific price P_n. House purchases may be financed with mortgage debt, as explained below. Housing is traded at ages $j = 1$ and $j = 2$, and all housing is sold at age $j = 3$. At age $j = 2$, households decide whether to leave their initial neighborhood and purchase housing elsewhere. Movers face transaction costs proportional to the value of their current house κP_n. The stock of houses in each neighborhood H_n is supplied inelastically.

We think of neighborhoods as school attendance zones within the same school district. Each neighborhood n is associated with a local school attended by all children in the neighborhood. The quality of a local school Q_n is a function of the local house price P_n relative to prices in other school zones. Specifically:

$$Q_n = \left(\frac{P_n}{\overline{P}}\right)^{\alpha} \tag{2}$$

where $\overline{P}$ is the population-weighted average house price, and α is the elasticity of school quality to relative house prices. As with school value added in our empirical work, Q_n represents the marginal benefit to a child of one year of education at the local school n.

While Equation (2) embeds our primary empirical result from Section 2, in the model we do not take a stand on precisely how house price changes are transmitted to local school quality. Our empirical results indicated that peer composition did not directly influence school quality, but that instead, high house price growth changed peer composition, which drew in higher quality teachers. In our model, house price differences across neighborhoods lead to different neighborhood composition: in this way school quality is semi-endogenously determined by local demographic factors.[14]

Human Capital, Test Scores, Productivity, and Income: Children are born with an innate ability level a_k, which is imperfectly inherited from their parents. Ability follows a log-AR(1) process:

$$\ln a_k = (1 - \rho_a)\mu_a + \rho_a \ln a_p + \varepsilon_a \tag{3}$$

where a_p is the ability of a parent, μ_a is average log-ability, ρ_a governs intergenerational persistence, and ε_a is an IID normal shock with mean zero and standard deviation σ_a.

Human capital is developed in childhood and is a function of underlying ability and

[14] See Fogli et al. (2025) and Chyn et al. (2025) for an example of a model with fully endogenous neighborhood quality due to neighborhood sorting on child ability and parent incomes. Zheng et al. (2021) model district-level school quality as a function of local tax revenues. Here, differences in quality across school zones is independent of local property tax revenues since these are received and allocated by the broader school district.

school quality:

$$h_k = a_k^{1-\gamma}(\tau Q_n)^\gamma \tag{4}$$

where τ is the amount of time spent in school and γ is the importance of school quality relative to child ability in human capital production. Because schools vary in educational quality Q_n, parents can directly influence human capital accumulation via their choice of neighborhood and thus the school their child attends.[15]

To map the model to the data, we assume that student test scores are measurable in the model. These do not affect child outcomes directly, but are a function of relative human capital accumulation. Test scores t_k are given by

$$t_k = \frac{h_k - \mu(h)}{\sigma(h)} \tag{5}$$

where $\mu(h)$ and $\sigma(h)$ are the population mean and standard deviation of human capital h_k. That is, test scores are normalized to have zero mean and unit standard deviation.

A child's earnings or labor productivity y_k is known upon entry to the labor force at age $j = 1$. Labor productivity is given by a Cobb-Douglas function of the child's human capital and their parent's labor productivity y_p:

$$y_k = h_k^{1-\delta} y_p^\delta \tag{6}$$

where δ governs the influence of parental productivity relative to accumulated human capital.[16] Between ages $j = 1$ and $j = 2$, labor productivity follows a log-random walk:

$$\ln y_2 = \ln y_k + \varepsilon_y$$

where ε_y is IID normal with standard deviation σ_y and mean $\mu_y = -\frac{1}{2}\sigma_y^2$. For tractability we assume no further shocks between ages $j = 2$ and $j = 3$, so $y_3 = y_2$.

Finally, income at each age is given by productivity and a deterministic, age-specific factor χ_j that captures the life-cycle profile of income. Thus, age j income is $\chi_j y_j$.

[15] For tractability we abstract from direct parent investments in child education such as parental time or other resources (see Cunha et al., 2010). Allowing for direct investments will dampen intergenerational inequality to the extent that investments are substitutes for school quality (Greaves et al., 2023), but will amplify intergenerational inequality if investments are complementary with school quality (Attanasio et al., 2022) or if financial constraints are more binding on poor households (Daruich et al., 2020).

[16] A long literature documents the influence of family income on educational attainment, independent of underlying ability. See for example Carneiro et al. (2002), Belley et al. (2007), and Leukhina (2023).

3.2. Household Decision Problems

At each age j, the household state vector $\{b, y, a, n\}$ consists of savings or debt b, labor productivity y, ability a, and current neighborhood n.

Decision Problem for Young Adults: At age $j = 1$, young adults are endowed with transfers provided by their parents b, labor productivity y, and their own ability a. Young adults do not have children or own housing. Their decision problem is divided into two sub-periods. First, idiosyncratic taste shocks over neighborhoods are realized and young adults choose a neighborhood to live in. Second, given a choice of neighborhood young adults consume c_1, and borrow or save b'_1.

The second sub-problem given a neighborhood choice n' is:

$$
\begin{aligned}
V_1(b, y, a; n') = \max_{c_1, b'_1} \; & \log(c_1) + \beta \mathbb{E}\left[V_2(b'_1, y', a_k, n')\right] \\
\text{s.t.} \quad & c_1 + b'_1 + P_{n'} = \chi_1 y + b \\
& \ln a_k = (1 - \rho_a)\mu_a + \rho_a \ln a + \varepsilon_a \\
& \ln y' = \ln y + \varepsilon_y \\
& b'_1 \geq -\theta P_{n'}
\end{aligned} \tag{7}
$$

where there is uncertainty over future child ability a_k and labor productivity y'. When $b'_1 < 0$ the household uses a mortgage to finance the house purchase. Borrowing is subject to a loan-to-value constraint, where θ is the maximum loan-to-value ratio. Expectations over future values $\mathbb{E}[V_2]$ are taken with respect to the evolution of child ability a_k, productivity shocks ε_y, and idiosyncratic neighborhood taste shocks $\varepsilon_{n'}$ arriving at age $j = 2$. Future values are discounted at rate β.

In the first sub-problem households choose a neighborhood n'_1 given common preference shifters Z_n and idiosyncratic taste shocks ε_n, both of which are unrelated to housing costs and school quality. Common preferences capture differences across neighborhoods valued by all households, such as the quality of housing stock or local amenities, while taste shocks reflect differences valued by individual households. The choice problem is:

$$
V_1(b, y, a) = \max_{n'_1} \left\{ V_1(b, y, a; n'_1) + Z_{n'_1} + \sigma_n \varepsilon_{n'_1} \right\} \tag{8}
$$

where Z_n are fixed and ε_n are drawn from a Type 1 Extreme Value distribution with scale parameter σ_n.

Decision Problem for Middle-Aged Adults: The age $j = 2$ problem is also divided into two sub-periods. First, children are born and their ability is revealed, idiosyncratic taste shocks over neighborhoods are realized, and adults may choose a new neighborhood to live in. Second, conditional on choice of neighborhood, parents consume c_2, borrow or save b'_2, and leave transfers for their children b'_k.

The second household sub-problem given neighborhood choice n' is:

$$
\begin{aligned}
V_2(b, y, a_k, n; n') = \max_{c_2, b_2', b_k'} \; & \log(c_2) + \beta V_3(b_2', y, n') + \varphi \log(b_k' + Y_k) \\
\text{s.t.} \quad & c_2 + b_2' + b_k' = \chi_2 y + (1+r)b + \mathbb{1}_{n' \neq n} (P_n - P_{n'} - \kappa P_n) \\
& h_k = a_k^{1-\gamma} (\tau Q_n)^{\gamma} \\
& y_k = h_k^{1-\delta} y_p^{\delta} \\
& Y_k = \chi_1 y_k + \frac{\chi_2 y_k}{1+r} + \frac{\chi_3 y_k}{(1+r)^2} \\
& b_2' \geq -\theta P_{n'}, \; b_k' \geq 0
\end{aligned} \tag{9}
$$

where b_2' is the choice of savings or debt, b_k' are transfers to children, h_k is the child's human capital, y_k is the child's labor productivity, and Y_k is the present value of the child's life-time income discounted at the interest rate r. If moving across neighborhoods, the household receives the proceeds from selling its old house and buying a new house $P_n - P_{n'}$ less the housing transaction cost κP_n. The household can also borrow $b_2' < 0$ subject to the loan-to-value constraint.

We assume that parents care about the life-time wealth of their children. This includes transfers b_k' and the present value of life-time income Y_k. We make two assumptions regarding parental transfers for tractability. First, parents ignore uncertainty over child income and focus only on the permanent component of labor productivity y_k.[17] Second, we assume that as in Fogli et al. (2025), parents value child outcomes via the same log-utility function over their own consumption, and the parameter φ governs the strength of their altruism. [18]

In the first sub-problem households choose a neighborhood n_2' subject to common preference shifters Z_n, idiosyncratic taste shocks ε_n, and a utility cost of moving across neighborhoods η. This cost is necessary for calibration, as households in the data are much less mobile than a frictionless model would imply.[19] The choice problem is:

$$
V_2(b, y, a_k, n) = \max_{n_2'} \left\{ V_2(b, y, a_k, n; n_2') + Z_{n_2'} + \sigma_n \varepsilon_{n_2'} - \eta \mathbb{1}_{n_2' \neq n} \right\} \tag{10}
$$

Decision Problem for Old Adults: Age $j = 3$ households consume income, the

[17] Our assumptions preclude the possibility of a dynastic precautionary savings mechanism as, for example, discussed by Boar (2021).

[18] These assumptions significantly simplify our computations as the model only needs to be solved backwards from age $j = 3$ once. That is, we do not need to recursively iterate over the solutions to parent and child value functions. Chyn et al. (2025) model parents as caring for children explicitly by including the child's continuation value in the parent's own value function. However, they simplify by modeling two neighborhoods whereas we have five.

[19] See, for example, Nakamura et al. (2022) who argue that large frictions are necessary to explain very low mobility rates even in the face of large monetary and family benefits from moving across location.

proceeds from selling their house, and any remaining assets. The value function is

$$
\begin{aligned}
V_3(b, y, n) = \; & \log(c_3) \\
\text{s.t.} \quad & c_3 = \chi_3 y + (1+r)b + P_n
\end{aligned}
\tag{11}
$$

3.3. Equilibrium

A stationary equilibrium consists of house prices $\{P_n\}$, decision rules for consumption $\{c_1, c_2, c_3\}$, borrowing or saving $\{b'_1, b'_2\}$, transfers $\{b'_k\}$, neighborhood choices $\{n'_1, n'_2\}$, and invariant distributions $\{\lambda_1, \lambda_2, \lambda_3\}$, such that: (i) given house prices, the decision rules solve the household problems (7)–(11); (ii) housing markets clear in each neighborhood

$$
\sum_{j=1,2} \int \mathbb{1}_{j,n} d\lambda_j = H_n \tag{12}
$$

(iii) school quality in each neighborhood is determined by Equation (2); (iv) and the stationary distributions satisfy

$$
\lambda_1 = \int Q_{2,K} d\lambda_2, \quad \lambda_2 = \int Q_{1,2} d\lambda_1, \quad \lambda_3 = \int Q_{2,3} d\lambda_2
$$

where $Q_{j,j'}$ are distribution transition functions from age j to j', and $Q_{2,K}$ is the transition function from parents at age $j = 2$ to children at age $j = 1$. For computational details, see Appendix B.

3.4. Calibration

The aim of our calibration strategy is to capture the cross-sectional inequality observed within our school district of interest. We first fix a number of parameters according to external information. We then use a moment matching exercise to set a number of internally calibrated parameters. We target a range of statistics reflecting cross-neighborhood dispersion in house prices, incomes, student test scores, and intergenerational mobility, along with statistics on intergenerational transfers and residential mobility. We primarily draw on statistics for the period 2000–2006, prior to the Financial Crisis.

Many of our statistics depend on mapping school zones in the data to model-consistent neighborhoods. To do this, we first group school zones in the data by house price, and then calculate population-weighted statistics within each group of zones. We compute median house prices for each school zone using ZTRAX data from 2000–2006 (Zillow, 2020) and group zones according to population-weighted quintiles of the house price distribution. Since many statistics of interest are not reported for school zones, we use a cross-walk between year 2000 census tract and school zone definitions.[20] We aggregate census tract statistics to the school zone level by computing population-weighted averages. Finally, we allocate school zone-level statistics to model neighborhoods according to the house price quintiles computed above.[21]

Because our administrative data on school test scores is not linked to information on children's future income, we must draw on other sources. For example, we make extensive use of intergenerational income mobility statistics provided by Chetty et al. (2018). These statistics are reported conditional on parent quantiles of the U.S. income distribution. However, our model represents only households in our school district of interest. To address this discrepancy, we map moments from the national income distribution into the district income distribution using the American Community Survey. Thus, all of our statistics are reported relative to district-level incomes. See Section B.3 in the Appendix for details.

Externally calibrated parameters: A model period is 15 years, model ages $j = 1, 2, 3$ represent households aged 21–35, 36–50, and 51–65, and the population size of each cohort is normalized to one. The supply of housing is fixed and neighborhood sizes are equal: $H_n = \frac{1}{5}$ for all n. Neighborhoods are distinguished by their house prices such that $P_A < P_B < P_C < P_D < P_E$. In the initial steady state equilibrium we normalize average prices so that $\overline{P} = \sum_n H_n P_n = 1$.

Panel (a) of Table 4 reports externally calibrated model parameters. The real annual interest rate is set to 2 percent, and for simplicity we assume the same interest rate on savings and mortgages. In the data, the median LTV ratio at origination is 80 percent. Since a typical mortgage maturity is 30 years and one model period is 15 years, we assume that households repay half of their mortgage principal within a model period. Hence, we set the maximum LTV ratio θ to 0.4. We set the housing transaction cost κ to 5 percent of the value of a house, consistent with data on real estate agent commissions in Banerjee et al. (2025). A key parameter in the model is the elasticity of school quality to local house price α from Equation (2). We take this parameter directly from our empirical estimates in Table 1, and set $\alpha = 0.25$. Given the number of years that children spend in elementary school, we set $\tau = 6$.

We calibrate β consistent with external information on marginal propensities to consume (MPCs). Berger et al. (2018) derive a sufficient statistic for MPCs out of permanent housing wealth shocks in life-cycle models with housing. This is given by the MPC out of

[20] See Appendix A.5 for details.

[21] To get a sense of the spatial concentration of school zones within neighborhoods, we calculate the fraction of a school zone's edges that border another zone assigned to the same neighborhood category. Under random assignment, this fraction would be about 20%. Neighborhood A exhibits substantial clustering, with 77% of its school-zone edges bordering another school zone belonging to the same neighborhood. In contrast, neighborhoods B, C, and D display limited geographic concentration, with only about 30–35% of their edges adjoining neighborhoods of the same type. The highest-priced neighborhood, "E" falls in between but is much closer to the less-clustered neighborhoods.

Table 4: Model Parameters

Description	Parameter	Value	Source
Panel (a): Externally Calibrated Parameters			
Annual interest rate	r	0.020	Authors
Maximum LTV ratio	θ	0.400	Authors
Moving cost	κ	0.050	Banerjee et al. (2025)
House price elasticity of school quality	α	0.250	Authors
Time spent in school	τ	6.000	Authors
Annual discount factor	$\beta^{\frac{1}{15}}$	0.850	See text
Life-cycle income profile	$\{\chi_j\}$	{1.00,1.42,1.51}	ACS, 2005–2009
Average ability	μ_a	0.202	Normalization
Panel (b): Internally Calibrated Parameters			
Altruism	φ	2.656	Calibrated
Utility cost of moving	η	0.731	Calibrated
Std. dev. neighborhood taste shocks	σ_n	0.397	Calibrated
Std. dev. ability shocks	σ_a	0.927	Calibrated
Intergenerational persistence of ability	ρ_a	0.396	Calibrated
Std. dev. income shocks	σ_y	0.771	Calibrated
Human capital elasticity of school quality	γ	0.566	Calibrated
Child income elasticity of parent income	δ	0.114	Calibrated
Neighborhood demand shifter, B	Z_B	0.252	Calibrated
Neighborhood demand shifter, C	Z_C	0.420	Calibrated
Neighborhood demand shifter, D	Z_D	0.643	Calibrated
Neighborhood demand shifter, E	Z_E	1.046	Calibrated

transitory income multiplied by the level of housing wealth. In a broad class of models, MPCs out of transitory income shocks are approximately given by $1 - \beta$.[22] So we set the annual discount factor to 0.85, consistent with estimated MPCs of around 0.15 for high-liquidity households (see Baugh et al., 2021; Graham et al., forthcoming). In Section 3.5 we show that the model generates empirically plausible MPCs out of housing wealth shocks.

The life-cycle profile of income $\{\chi_j\}_{j=1,2,3}$ is computed from the ratios of average incomes between ages 36–50 and 51–65 relative to average incomes between ages 21–35 using data from the American Community Survey (ACS) 5-year data for 2005–2009 (U.S. Census Bureau, 2019). Conditional on all other parameters, we normalize the mean of the ability process μ_a to ensure that the lowest income household at age $j = 1$ can afford to purchase a house in the cheapest neighborhood A.[23] See Appendix B.2 for details.

Internally calibrated parameters: Panel (b) of Table 4 reports model parameters chosen via a moment matching exercise. We set the parameters $\{\varphi, \eta, \sigma_n, \sigma_a, \rho_a, \sigma_y, \gamma, \delta, Z_B, Z_C, Z_D, Z_E\}$ to target the statistics reported in Table 5. Although the parameters are calibrated jointly, some statistics are more informative about

[22] In a finite-horizon permanent income hypothesis model, the MPC is $\frac{1-\beta}{1-\beta^J}$ where J is remaining years of life and the MPC converges to $1 - \beta$ for sufficiently large J.

[23] Similarly, Fogli et al. (2025) model all households as renters and normalize rent in the cheapest neighborhood to zero.

Table 5: Moments used in Model Calibration

Moment	Model	Data	Source
Transfers share of networth	0.253	0.260	Feiveson et al. (2018)
Average income ratio, E/A	2.479	2.760	ACS, 2005–2009
Income p_{90}/p_{10}	9.878	9.710	ACS, 2005–2009
Average rank(child income$\|p_{25}$), $E-A$	0.091	0.088	Chetty et al. (2018)
Average rank(child income$\|p_{75}$), $E-A$	0.075	0.080	Chetty et al. (2018)
Move probability, $j=2$	0.346	0.441	CPS, 2000–2006
Moving regression R^2	0.167	0.194	PSID, 2001–2019
Average test score difference, $E-A$	0.835	0.837	Admin. data, 2003–2006
House price ratio, B/A	1.370	1.370	Zillow, 2000–2006
House price ratio, C/A	1.680	1.680	Zillow, 2000–2006
House price ratio, D/A	2.170	2.170	Zillow, 2000–2006
House price ratio, E/A	3.330	3.330	Zillow, 2000–2006

certain parameters than others. Appendix B.5 presents a parameter sensitivity analysis, which informs parameter identification and the discussion that follows.

We set the altruism parameter φ to target the aggregate ratio of life-time within-family transfers to networth, as reported by Feiveson et al. (2018). The utility cost of moving η targets the probability of a cross-neighborhood move between ages $j=1$ and $j=2$. To calculate this statistic in the data we use the 2000–2006 waves of the Current Population Survey (Flood et al., 2023). Consistent with parent-households in the model, we restrict the sample to married homeowners aged 35–50, and define cross-neighborhood moves as within-county moves over the last year. Assuming that households move at most once between ages 35 and 50, the probability of moving at any time in this 15-year period is: $\pi_{35}+\sum_{j=36}^{50}\pi_j\times\prod_{s=35}^{j-1}(1-\pi_s)$.

We set the standard deviation of neighborhood taste shocks σ_n to target the R^2 statistic from a regression predicting residential mobility. This captures the extent to which cross-neighborhood moves are predicted by household state variables observed in the model. A low R^2 suggests that moves are not well-explained by household circumstances, and are thus driven by dispersion in idiosyncratic neighborhood preference shocks. Using data from the Panel Study of Income Dynamics for 2001–2019 (Social Research Center, 2019), we regress an indicator variable for whether a household moved on their age, their lagged log-household income, and their lagged log-house price. Consistent with our model, we restrict the sample to homeowners only. Appendix B.4 provides additional details.

We choose the dispersion of labor productivity shocks σ_y to target the ratio of household incomes at the 90^{th} percentile relative to the 10^{th} percentile. We then set the standard deviation of ability shocks σ_a to target the ratio of average incomes in neighborhood E relative to neighborhood A. The larger is dispersion in innate ability the stronger is sorting across neighborhoods by child ability and parent income, since these tend to be correlated. In the data we compute incomes across the population and across neighborhoods using the 2005–2009 ACS (U.S. Census Bureau, 2019).

We set the intergenerational persistence of ability ρ_a and the elasticity of child to parent labor productivity δ to jointly target intergenerational income mobility across our neighborhoods. Chetty et al. (2018) calculates various intergenerational mobility statistics for children born in the 1980s with subsequent earnings measured in 2011–2012. Our preferred measure is the average income rank of children, since this captures the possibility of both upward and downward mobility conditional on parent income and neighborhood. We use the average income ranks of children born to parents at the 25^{th} and 75^{th} percentiles of the income distribution, and we target the differences in these rank measures between neighborhoods E and A.[24]

The human capital elasticity of school quality γ is chosen to target differences in average school test scores across neighborhoods. We compute average test scores by neighborhood in our administrative data for 2003–2006. We target the difference in average normalized test scores between neighborhoods E and A.

Finally, the neighborhood preference shifter Z_A is normalized to zero, and we set the remaining preference shifters Z_n to target house prices relative to neighborhood A: P_n/P_A for $n \in \{B, C, D, E\}$. We take house price data from ZTRAX for 2000–2006 (Zillow, 2020).

Model Fit: Following our calibration, we evaluate model fit along two dimensions.

First, Figure 2 illustrates a range of cross-neighborhood statistics in the model and data. Panel (a) illustrates average neighborhood incomes relative to A and shows that incomes rise consistently from the lowest- to highest-price neighborhoods (U.S. Census Bureau, 2019). In the model this reflects the strong sorting of high income households into neighborhoods with higher quality schools and higher common preference shifters Z_n. Panel (b) illustrates average student test scores reported relative to the middle neighborhood C. Although our calibration only targets the differences in average test scores between E and A, we find that test scores rise monotonically with neighborhood house prices as in the data. In the model this is due to the direct effects of higher school quality as well as endogenous neighborhood sorting according to child ability. Panel (c) illustrates the probability that a child leaves the neighborhood that their parents raised them in by the time that they are adults (Chetty et al., 2018). While this is an untargeted moment in our calibration, we find that in both model and data the average probability is around 75 percent. In the data the probability is consistently rising with neighborhood house price, while in the model the pattern is hump-shaped across the house price distribution.

[24]As noted above, we take parents at the 25^{th} and 75^{th} percentiles of the school district's income distribution. See Appendix B.3 for details of the mapping between the national- and district-level income distributions.

Figure 2: Cross-Neighborhood Statistics in Model and Data

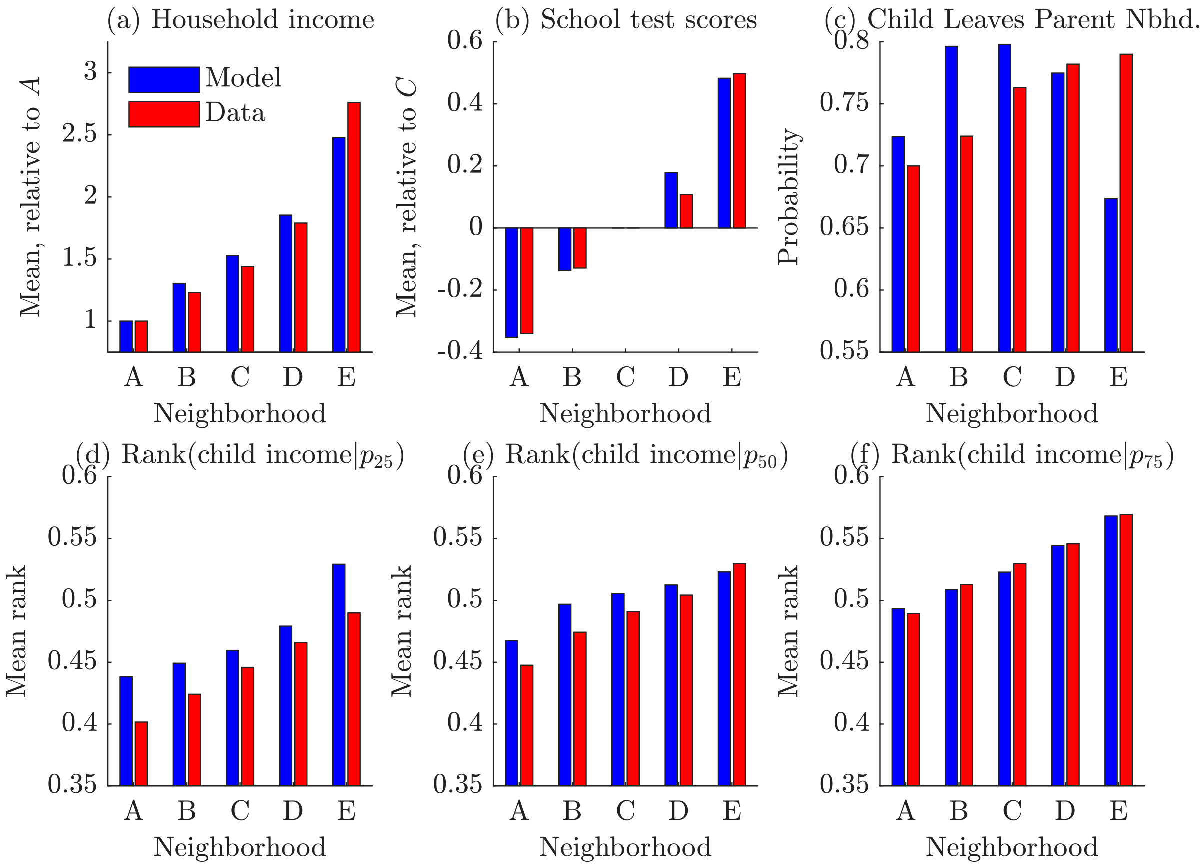

Source: Authors' calculations using administrative data, U.S. Census Bureau (2019), Chetty et al. (2018).

Panels (d)–(f) report average child income ranks for parents at the 25^{th}, 50^{th}, and 75^{th} percentiles of the income distribution (Chetty et al., 2018). While our calibration only targets the differences in average income ranks for parents at the 25^{th} and 75^{th} percentiles and between neighborhoods E and A, the figures show that our model matches the levels and cross-neighborhood dispersion of intergenerational income mobility well. We can see that children with higher income parents have higher average income ranks, while average child income rank is increasing with neighborhood house price regardless of parent income. That is, better neighborhood schools can improve outcomes for children growing up in families at all points in the income distribution.

Second, we evaluate the model against estimates of the intergenerational effects of the Moving to Opportunity (MTO) experiment studied by Chetty et al. (2016).[25] The MTO experiment randomized the distribution of housing vouchers to families in Baltimore, Boston, Chicago, Los Angeles, and New York. Families were eligible for a voucher if their income was below 50 percent of median household income and if they lived in public housing in low-income neighborhoods. Vouchers ensured that families paid no more than

[25] Chyn et al. (2025) and Gregory et al. (2022) conduct similar validation exercises in their models.

30 percent of their labor income on rent, but they could only be used for private market housing in neighborhoods with poverty rates of less than 10 percent. Chetty et al. (2016) finds that 48 percent of eligible families accepted a voucher, and they estimate that the children of parents who moved earned 31 percent more as adults relative to the control group.

We run a similar experiment in the model by offering a voucher to age $j = 2$ households living in neighborhood A with income below 50 percent of median income. Households can use the voucher to move to any of neighborhoods $\{B, C, D, E\}$, and only pay $\min\{P_n, 0.3 \times \textit{income}\}$ for their housing. The experiment is conducted in partial equilibrium, with no associated change in house prices or school qualities. We find that 99.5 percent of eligible households accept the voucher, 86 percent of households who would not otherwise move are induced to do so, and 85 percent of voucher recipients move to the most expensive neighborhood E. Computing the treatment effect on the treated, we find that the voucher causes a 14 percent increase in average future child incomes. Note that the structure and calibration of our model imply that the maximum possible increase in child incomes due to moving from neighborhood A to E is 16 percent.[26] We would expect larger effects in a model with additional channels of child human capital investment and where there is complementarity between investment and school quality.[27]

3.5. Intra- and Intergenerational Wealth Effects

We now conduct model-based experiments to quantify the wealth effects of housing market shocks through the school quality channel. We run five dynamic simulations, and in each simulation one of the five neighborhoods is hit by a shock. The shocks are modeled as unexpected, permanent, common neighborhood demand shocks ΔZ_n. We first solve for new steady state equilibria under each shock, where ΔZ_n is calibrated so that the long-run equilibrium price in neighborhood n increases by 10 percent. Remaining neighborhood prices adjust to clear their respective housing markets, and school quality Q_n in each neighborhood adjusts according to Equation (2). We then solve for dynamic equilibrium transition paths to the new steady state: $\{P_{n,t}, Q_{n,t}\}$.

As an example of one of our simulations, Figure 3 illustrates the transition paths following a demand shock to neighborhood C.[28] Following the shock, demand to live in C rises relative to other neighborhoods, which leads to a rise in the equilibrium house price P_C. A higher price also increases school quality Q_C relative to other neighborhoods, which further increases demand among parent households. Additionally, households re-sort across neighborhoods so that average student ability in each neighborhood also adjusts.

[26] From Equations (4) and (6) we have $y_k(Q_E)/y_k(Q_A) = (Q_E/Q_A)^{\gamma(1-\delta)}$.

[27] See for example, Fogli et al. (2025), Attanasio et al. (2022), Gregory et al. (2022), and Chyn et al. (2025).

[28] See Appendix B.6 for the transition paths for all 5 simulated experiments.

Figure 3: Transition Paths Following Demand Shock to Neighborhood C

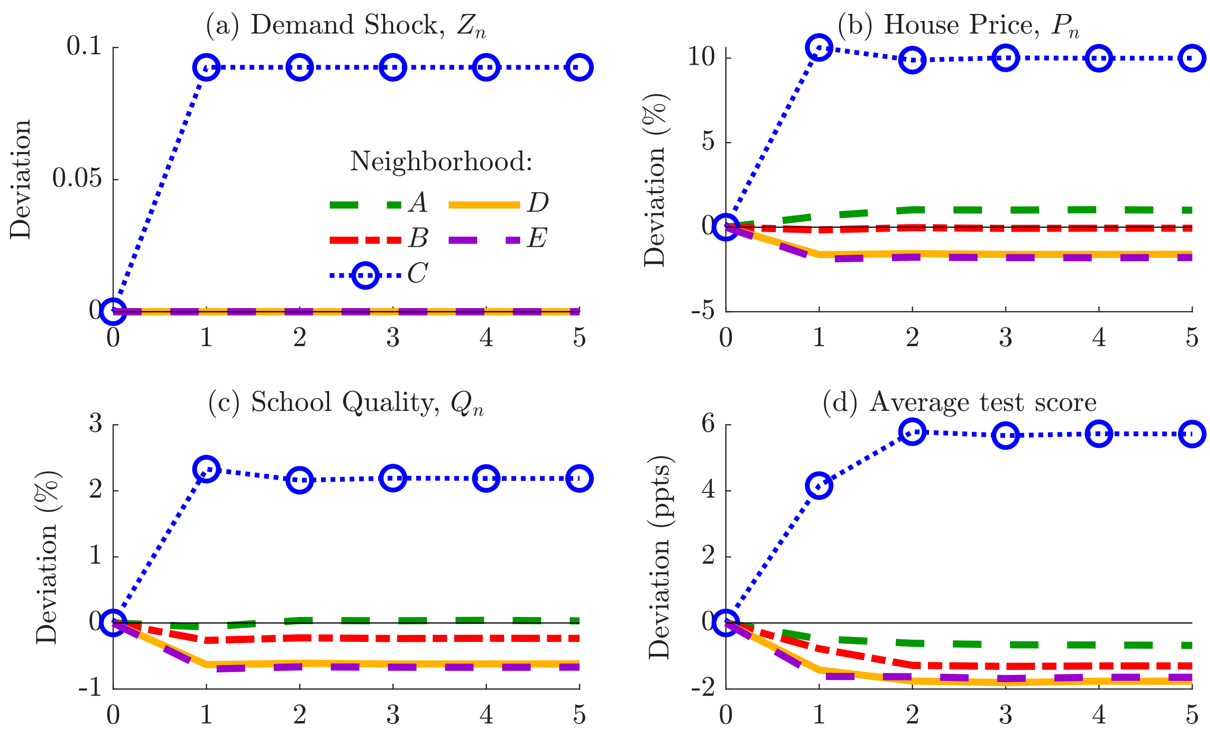

As a result, average test scores in C also improve but more slowly than what is implied by the adjustment in school quality alone. While other neighborhoods are not directly impacted by the demand shock, their house prices and school quality adjust as part of the spatial equilibrium.

We now quantify the intra- and intergenerational wealth effects of housing market shocks. We keep track of age $j = 2$ households in the first period $t = 1$ along the equilibrium transition path following a shock. These households decide how much to consume, save, transfer to children, and whether to move neighborhood. The responses of consumption and transfers to the shocks reflect intra- and intergenerational wealth effects, respectively. The choice to stay or move across neighborhoods affects a child's school quality and future income, reflecting an additional intergenerational wealth effect.

To compute wealth effects, we simulate a sample of 25,000 age $j = 2$ households and record decisions and outcomes with and without the housing market shocks. A household has the same state vector $\{b, y, a_k, n\}$ in steady state and at the time of the shock before making decisions. We estimate the marginal effect β_x of house prices on decision or outcome x via the regression

$$\Delta x_{i,n,t} = \beta_x \Delta P_{n,t} + \alpha_N + \epsilon_{i,n,t} \tag{13}$$

where i denotes a household, n is the household's initial neighborhood at the time of the shock, $\Delta x_{i,t} = x_{i,t} - x_i$ and $\Delta P_{n,t} = P_{n,t} - P_n$ are time-t deviations from steady state, and $\epsilon_{i,n,t}$ is an error term. We pool data from households across the five housing shock experiments, and a regression fixed effect α_N denotes the experiment with shock to neighborhood N. The outcome variables of interest are consumption c, transfers b'_k, the net present value of child life-time income Y_k, and the probability of moving to a new neighborhood $\mathbb{P}(n' \neq n)$. Finally, we report all wealth effects at an annual frequency by dividing by the 15-year model period.

Table 6: Housing Wealth Effects of Neighborhood Demand Shocks

	Marginal propensity to consume	Marginal propensity to transfer	Marginal change in child income	Marginal propensity to move
Panel (a): All households				
All	0.024	0.012	0.017	-0.015
Panel (b): Heterogeneous responses by demographics				
Low child ability	0.018	0.017	0.011	-0.012
Low parent income	0.026	0.002	0.021	-0.021
Low parent networth	0.028	0.010	0.021	0.001
Panel (c): Heterogeneous responses by neighborhood				
Low initial price	0.028	0.010	0.021	0.000
Price decrease	0.030	0.016	0.008	-0.011

Notes: Wealth effects estimated for age $j = 2$ households from Equation (13) and reported in annualized terms. Heterogeneous effects estimated by modifying Equation (13) to include interactions terms with price growth $\Delta P_{n,t}$. Low ability, income, networth, and initial house price are defined as households in lowest quintile of the distribution of values for each variable, respectively.

Table 6 reports our results. Panel (a) shows average annualized marginal effects across all households. The marginal propensity to consume (MPC) out of house prices is 0.024 (that is, 2.4 cents in the dollar), the marginal propensity to transfer housing wealth to children is 0.012 (1.2 cents in the dollar), and the marginal change in life-time income for children is 0.017 (1.7 cents in the dollar). Additionally, the marginal propensity to move neighborhoods is -0.015. Positive house price shocks increase housing wealth for households who then consume a little more and transfer a little more to their children. Households are also less likely to move neighborhoods following the shock, which implies that they are more likely to take advantage of increasing local school quality. The effects of higher school quality then contribute to higher future incomes for children – an indirect intergenerational wealth effect. We discuss heterogeneous effects in Section 3.6.

Our model-based housing wealth effects are consistent with empirical estimates from the literature. Our average annualized MPC of 0.024 is similar to recent estimates for non-durable goods in the range of 0.01 to 0.03 (Mian et al., 2013; Aladangady, 2017; Graham et al., 2023). This literature also suggests that households with low initial networth, or higher housing leverage, experience larger wealth effects (see Panel (b) of Table 6). Marginal propensities to transfer out of housing wealth are less well-studied in the literature. However, recent empirical work by Daysal et al. (2023) estimates that

8 to 16 percent of housing wealth shocks experienced by parents are transmitted to the housing wealth of children.[29] Converting our marginal propensity to an elasticity, we find that 3 percent of housing wealth shocks are transmitted to children in the form of cash transfers.[30]

We evaluate the robustness of these results by computing wealth effects under a different set of housing market shocks. Specifically, we model unexpected, permanent, housing supply shocks ΔH_n. These represent long-run changes in the relative availability of housing across neighborhoods. As in our baseline experiments, we first solve for steady state equilibria under each shock, where ΔH_n is calibrated so that the long-run equilibrium price in neighborhood n increases by 10 percent. Table B.1 in Appendix B.7 reports our results. We find that housing wealth effects are similar under both neighborhood demand and supply shocks. Under the supply shocks, the MPC is 0.023, the marginal propensity to transfer to children is 0.011, and the marginal change in life-time income for children is 0.018. Only the marginal propensity to move is different, with a positive value of 0.011. This largely reflects the forced movement of many households as supply shrinks in some neighborhoods relative to others.

Table 7: Dollar-Valued Housing Wealth Effects

	Expenditures				
	Consumption	Transfers	Net housing	Saving	Child income
Dollar Values	\$46,257	\$24,002	\$-30,686	\$-39,759	\$33,924
Fraction of house price increase	0.36	0.18	-0.24	-0.31	0.26

Notes: Real, dollar-valued wealth effects derived from marginal wealth effects estimated via Equation (13), and evaluated at the empirical standard deviation of 5-year real house price changes. Marginal effects on consumption, transfers, and child income reported in Table 6.

In Table 7 we provide a simple summary of our main results by computing real, dollar-valued wealth effects given the empirical standard deviation of house price changes. For comparison with our empirical results in Section 2, we use the standard deviation of real 5-year growth rates (65 percent) multiplied by the median house price in our school district in the year 2000 (\$200,000). This yields a \$130,000 increase in house prices. We multiply this house price change by (non-annualized) marginal wealth effects calculated via Equation (13) (see Table 6 for annualized values). We report dollar-valued housing wealth effects for consumption, transfers to children, net housing expenditures, saving (or borrowing), and future child incomes. Note that consumption, transfers, net housing costs, and saving are all expenditures that appear in the age $j = 2$ household budget

[29] Benetton et al. (2022) find that parental home equity extraction out of housing wealth gains is associated with an increased probability that adult children transition to homeownership.

[30] We calculate the elasticity as $MPT \times \frac{mean(P_{n,SS})}{mean(b'_{k,SS})}$, taking the mean of house prices and transfers from the 25,000 simulated households in steady state.

constraint. In contrast, future child incomes are a function of school quality and involve no direct household spending.

From the $130,000 increase in house prices, we calculate that household consumption rises by $46,000, transfers to children rise by $24,000, net housing expenses fall by $31,000, and saving falls (or borrowing rises) by $40,000.[31] Additionally, the present discounted value of future child incomes rises by $34,000. The decline in net housing expenditure reflects both a decline in the likelihood of moving across neighborhoods (see Table 6), and that the net costs of moving are lower when houses in an initial neighborhood rise in value relative to others. The decline in saving or increase in borrowing is a result of consumption smoothing. Since many households do not move residence, the only way to access the increase in their home equity is to borrow and repay at age $j = 3$ when they finally sell their house.

These results show that intra-generational wealth effects through contemporaneous consumption represent 36 percent of the overall increase in house prices, while the inter-generational wealth effect through cash transfers to children makes up another 18 percent of the increase in house prices. An additional intergenerational wealth effect occurring through the increase in future children's incomes represents a further 26 percent of the increase in house prices. In sum, the intergenerational effect through children's incomes accounts for one-third of the total housing wealth effects at the time of a shock.

We can now compare our model-based wealth effects to our empirical results. Again, our model results suggest that a one standard deviation increase in local house prices is associated with a $34,000 increase in life-time income for a child. In Section 2.4 our back-of-the-envelope calculation suggests that a one standard deviation increase in local house prices is associated with a $37,000 increase in life-time incomes for children. Note the striking similarity in these estimates despite the fact that no feature of our model is calibrated to match the dynamic interaction between house prices and future incomes for children.

3.6. Additional Channels of Housing Wealth Effects

Finally, we use the model to explore additional transmission channels of housing market shocks. First, in Panels (b) and (c) of Table 6 we consider heterogeneity in housing wealth effects according to household and neighborhood characteristics.[32] We find that families with low-ability children have smaller MPCs, larger propensities to transfer, and smaller changes in future child incomes. Low ability children face lower returns to school quality, and so benefit less from improvements in their local school. As a result, their

[31] Note that age $j = 3$ consumption also rises when the house is sold in the following period.

[32] We estimate these heterogeneous effects by adding interaction terms to Equation (13). For conditioning variable w, we estimate $\Delta x_{i,n,t} = \beta_x \Delta P_{n,t} + \gamma_{x,w} \Delta P_{n,t} \times \mathbb{1}_w + \alpha_w \mathbb{1}_{i,w} + \alpha_N + \epsilon_{i,n,t}$ where $\mathbb{1}_{i,w}$ is an indicator variable for households with characteristic w.

Figure 4: Changes in Intergenerational Income Mobility

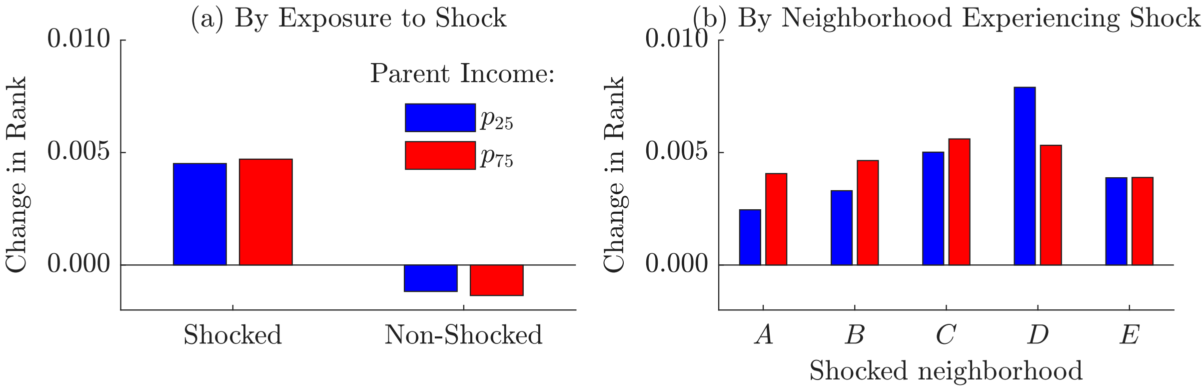

Notes: Intergenerational income mobility measured via average child income rank for parents at the 25^{th} and 75^{th} percentiles of the income distribution.

parents compensate them more by providing larger transfers out of housing wealth gains. Families with low-income parents have larger MPCs, smaller propensities to transfer, larger increases in future child incomes, and larger declines in residential movement. Low-income households have higher marginal utilities of consumption, hence large MPCs. They also choose to stay in place and take advantage of higher local school quality, which substitutes for the much lower transfers they provide to their children. Families with low-networth parents and those in low-house price neighborhoods behave similarly to families with low income, except for near-zero marginal propensities to move. Families experiencing house price decreases exhibit larger MPCs, larger marginal propensities to transfer wealth, smaller changes in future child incomes, and smaller declines in the propensity to move. This suggests that families hit with negative housing shocks largely bear the cost via lower contemporaneous consumption and transfers, but manage to avoid especially large declines in their children's future incomes.

Second, we show how these housing market shocks affect intergenerational income mobility. In Figure 4 we report changes in average child income ranks conditional on parent income. In Panel (a) we report averages of these changes across all neighborhoods experiencing positive housing demand shocks and across those that do not. Children in shocked neighborhoods experience substantial increases in income mobility. On average, the increase due to the shock is nearly 10 percent of the rank-difference between the top and bottom neighborhoods in steady state (see Table 5). Panel (b) shows changes in average income ranks conditional on parents being in the particular neighborhood experiencing the shock. Children in all neighborhoods experience increases in their income rank, however these gains are generally increasing with initial neighborhood house price.

Finally, we further investigate residential mobility decisions following a housing market shock. While neighborhood choices are sticky in steady state due to both monetary and utility costs of moving, Table 6 shows reduced residential mobility in response to shocks.

Figure 5: Changes in Residential Moving Probabilities

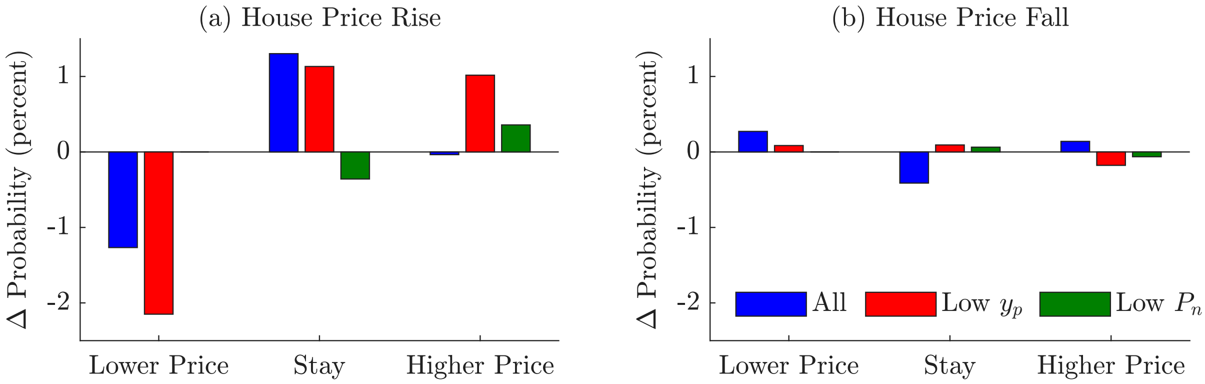

This reflects households' desire to stay in place to take advantage of improvements in local school quality. Figure 5 illustrates household moving decisions in more detail. Panel (a) shows that on average households experiencing positive house price shocks are much more likely to stay in their initial neighborhood, and much less likely to move to neighborhoods with lower prices and thus school qualities. Low income households are even less likely to move to lower priced neighborhoods, and are equally likely to stay in place or move to higher priced neighborhoods with better schools. Households from low-initial price neighborhoods are less likely to stay in place, and are slightly more likely to move to better neighborhoods. Panel (b) shows that households facing negative house price shocks, particularly low income households, are somewhat less likely to stay in place and are more likely to move to worse neighborhoods.

Overall, we find that housing market shocks have strong spatial and cross-sectional implications for intergenerational inequality. Positive house price shocks give households the opportunity to take advantage of improving local schools or to move up the neighborhood and school quality ladder. For those that stay, higher local school quality increases the income mobility of their children. Low-income families particularly benefit from these housing market shocks, as they can substantially increase both household consumption and children's future incomes. But the negative consequences of housing market shocks are also apparent. Families facing falling house prices see greater declines in consumption, a reduction in child incomes, and may be more likely to fall down the neighborhood and school quality ladder.

4. Conclusion

In this paper we study intra- and intergenerational housing wealth effects in the presence of a semi-endogenous school quality channel. First, using data from a large US school district we show that rising local house prices are associated with subsequent improvements in local school quality. Second, we quantify the wealth effects of house price shocks in an overlapping generations model with neighborhood choice, spatial equilibrium, and endogenous local school quality. Following an increase in house prices parents may consume more, transfer wealth to their children, and provide better educational opportunities for children through access to improving local schools. We find that intergenerational effects are large, with the effect of the school quality channel on children's future incomes accounting for one-third of total housing wealth effects.

Having documented this new channel for housing wealth effects, future research might consider the policy implications that follow. Importantly, policies need to account for our finding that the consequences of housing market shocks are broader than just the contemporaneous effects on current homeowners. For example, policies such as capital gain taxes would only target contemporaneous wealth gains and not the associated changes in intergenerational income mobility. Instead, policies may need to find a way to break the link between local house prices and local school quality. For example, school districts could allow students to attend public schools outside the school zone in which they live. Alternatively, schools and school districts could use counter-cyclical financial incentives to keep high-quality teachers in place following adverse local housing shocks.

References

Aladangady, Aditya (2017), "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata", *American Economic Review* 107, 3415–3446.

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

Allende, Claudia (2019), "Competition under social interactions and the design of education policies", *Job Market Paper*.

Allensworth, Elaine, Stephen Ponisciak, and Christopher Mazzeo (2009), "The schools teachers leave: Teacher mobility in Chicago public schools.", *Consortium on Chicago School Research*.

Altonji, Joseph G., Ching-I Huang, and Christopher R. Taber (2015), "Estimating the cream skimming effect of school choice", *Journal of Political Economy* 123, 266–324.

Arturo Santelli, Francisco and Jason A. Grissom (2022), *A Bad Commute: Does Travel Time to Work Predict Teacher and Leader Turnover and Other Workplace Outcomes?*, Working Paper, Annenberg Institute at Brown University.

Attanasio, Orazio, Teodora Boneva, and Christopher Rauh (2022), "Parental beliefs about returns to different types of investments in school children", *Journal of Human Resources* 57, 1789–1825.

Bacher-Hicks, Andrew, Thomas J. Kane, and Douglas O. Staiger (2014), *Validating Teacher Effect Estimates Using Changes in Teacher Assignments in Los Angeles*, Working Paper, National Bureau of Economic Research.

Banerjee, Rupkatha and Andrew Paciorek (2025), *Commissions and Omissions: Trends in Real Estate Broker Compensation*, FEDS Notes, Washington, D.C.: Board of Governors of the Federal Reserve System, DOI: 10.17016/2380-7172.3711, URL: https://www.federalreserve.gov/econres/notes/feds-notes/commissions-and-omissions-trends-in-real-estate-broker-compensation-20250512.html.

Bates, Michael D., Michael Dinerstein, Andrew C. Johnston, and Isaac Sorkin (2022), *Teacher Labor Market Equilibrium and the Distribution of Student Achievement*, Working Paper, National Bureau of Economic Research.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A Parker (2021), "Asymmetric consumption smoothing", *American Economic Review* 111, 192–230.

Baum-Snow, Nathaniel and Lu Han (2024), "The microgeography of housing supply", *Journal of Political Economy* 132, 1897–1946.

Belley, Philippe and Lance Lochner (2007), "The changing role of family income and ability in determining educational achievement", *Journal of Human capital* 1, 37–89.

Benabou, Roland (1994), "Human capital, inequality, and growth: A local perspective", *European Economic Review* 38, 817–826.

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

Benabou, Roland (1996b), "Heterogeneity, stratification, and growth: macroeconomic implications of community structure and school finance", *American Economic Review*, 584–609.

Benetton, Matteo, Marianna Kudlyak, and John Mondragon (2022), "Dynastic Home Equity", *Available at SSRN 4158773*.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra (2018), "House prices and consumer spending", *The Review of Economic Studies* 85, 1502–1542.

Black, Sandra E. (1999), "Do better schools matter? Parental valuation of elementary education", *Quarterly Journal of Economics* 114, 577–599.

Boar, Corina (2021), "Dynastic precautionary savings", *Review of Economic Studies* 88, 2735–2765.

Bonhomme, Stéphane, Grégory Jolivet, and Edwin Leuven (2016), "School characteristics and teacher turnover: Assessing the role of preferences and opportunities", *Economic Journal* 126, 1342–1371.

Boyd, Don, Hamp Lankford, Susanna Loeb, Matthew Ronfeldt, and Jim Wyckoff (2011), "The role of teacher quality in retention and hiring: Using applications to transfer to uncover preferences of teachers and schools", *Journal of Policy Analysis and Management* 30, 88–110.

Brandsaas, Eirik Eylands (2025), "Illiquid homeownership and the bank of mom and dad".

Brown, Catherine, Scott Sargrad, and Meg Benner (2017), "Hidden Money", *Centre for American Progress*.

Carneiro, Pedro and James J Heckman (2002), "The evidence on credit constraints in post-secondary schooling", *The Economic Journal* 112, 705–734.

Carver-Thomas, Desiree and Linda Darling-Hammond (2017), *Teacher Turnover: Why It Matters and What We Can Do About It*, tech. rep., Palo Alto, CA: Learning Policy Institute.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter (2018), *The opportunity atlas: Mapping the childhood roots of social mobility*, tech. rep., National Bureau of Economic Research.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff (2014a), "Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates", *American Economic Review* 104, 2593–2632.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff (2014b), "Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood", *American Economic Review* 104, 2633–2679.

Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz (2016), "The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment", *American Economic Review* 106, 855–902.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (2014c), "Where is the land of opportunity? The geography of intergenerational mobility in the United States", *Quarterly Journal of Economics* 129, 1553–1623.

Chyn, Eric and Diego Daruich (2025), "An equilibrium analysis of the effects of neighborhood-based interventions on children", *American Economic Review* 115, 4476–4522.

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

Daruich, Diego and Julian Kozlowski (2020), "Explaining intergenerational mobility: The role of fertility and family transfers", *Review of Economic Dynamics* 36, 220–245.

Davis, Matthew and Fernando Ferreira (2022), "Housing disease and public school finances", *Economics of Education Review* 88, 102236, ISSN: 0272-7757.

Daysal, N Meltem, Michael F Lovenheim, and David N Wasser (2023), *The intergenerational transmission of housing wealth*, tech. rep., National Bureau of Economic Research.

Deming, David J. (2014), "Using School Choice Lotteries to Test Measures of School Effectiveness", *American Economic Review* 104, 406–11.

Durlauf, Steven (1996a), "Neighborhood Feedbacks, Endogenous Stratification, and Income Inequality", *Dynamic Disequilibrium Modelling: Proceedings of the Ninth International Symposium on Economic Theory and Econometrics, W. Barnett, G. Gandolfo, and C. Hllinger, eds.*

Durlauf, Steven N. (1996b), "A theory of persistent income inequality", *Journal of Economic Growth* 1, 75–93.

Eckert, Fabian and Tatjana Kleineberg (2019), "Can we save the American dream? A dynamic general equilibrium analysis of the effects of school financing on local opportunities", *2019 Meeting Papers*, Society for Economic Dynamics.

Feiveson, Laura and John Sabelhaus (2018), "How does intergenerational wealth transmission affect wealth concentration?", *FEDS Notes. Washington: Board of Governors of the Federal Reserve System.*

Fernández, Raquel and Richard Rogerson (1996), "Income distribution, communities, and the quality of public education", *Quarterly Journal of Economics* 111, 135–164.

Fernández, Raquel and Richard Rogerson (1998), "Public education and income distribution: A dynamic quantitative evaluation of education-finance reform", *American Economic Review*, 813–833.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry (2023), *Integrated Public Use Microdata Series, Current Popula-*

tion Survey: Version 9.0 [dataset], Data retrieved from IPUMS http://doi.org/10.18128/D030.V10.0.

Fogli, Alessandra, Veronica Guerrieri, Mark Ponder, and Marta Prato (2025), *The End of the American Dream? Inequality and Segregation in US Cities*, tech. rep., Federal Reserve Bank of Minneapolis.

Gilraine, Michael and Nolan G. Pope (2021), *Making Teaching Last: Long-Run Value-Added*, Working Paper, National Bureau of Economic Research.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020), "Bartik instruments: What, when, why, and how", *American Economic Review* 110, 2586–2624.

Graham, James and Christos A Makridis (2023), "House prices and consumption: a new instrumental variables approach", *American Economic Journal: Macroeconomics* 15, 411–43.

Graham, James and Robert A McDowall (forthcoming), *Mental Accounts and Consumption Sensitivity Across the Distribution of Liquid Assets*.

Greaves, Ellen, Iftikhar Hussain, Birgitta Rabe, and Imran Rasul (Feb. 2023), "Parental Responses to Information about School Quality: Evidence from Linked Survey and Administrative Data", *Economic Journal*.

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

Gyourko, Joseph, Albert Saiz, and Anita Summers (2008), "A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index", *Urban studies* 45, 693–729.

Hansen, Michael and Diana Quintero (2017), *Scrutinizing equal pay for equal work among teachers*, tech. rep., Washington, D.C.: Brookings Institution.

Johnston, Andrew C. (2020), *Teacher Preferences, Working Conditions, and Compensation Structure*, Working Paper, IZA Institute of Labor Economics.

Karbownik, Krzysztof (2020), "The effects of student composition on teacher turnover: Evidence from an admission reform", *Economics of Education Review* 75, 101960.

Koedel, Cory and Jonah E. Rockoff (2015), "Value-added modeling: A review", *Economics of Education Review* 47, 180–195.

Kotera, Tomoaki and Ananth Seshadri (2017), "Educational policy and intergenerational mobility", *Review of Economic Dynamics* 25, 187–207.

Leukhina, Oksana (2023), "The Changing Role of Family Income in College Selection and Beyond.", *Review (00149187)* 105.

Mian, Atif, Kamalesh Rao, and Amir Sufi (2013), "Household Balance Sheets, Consumption, and the Economic Slump", *Quarterly Journal of Economics* 128.

Morris, Carl N. (1983), "Parametric empirical Bayes inference: Theory and applications", *Journal of the American Statistical Association* 78, 47–55.

Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson (2022), "The gift of moving: Intergenerational consequences of a mobility shock", *The Review of Economic Studies* 89, 1557–1592.

National Center for Education Statistics (2018), *School Attendance Boundary Survey 2015-2016*, https://nces.ed.gov/programs/edge/SABS, (accessed September 2018).

Rothstein, Jesse (2015), "Teacher quality policy when supply matters", *American Economic Review* 105, 100–130.

Rothstein, Jesse M. (2006), "Good principals or good peers? Parental valuation of school characteristics, Tiebout equilibrium, and the incentive effects of competition among jurisdictions", *American Economic Review* 96, 1333–1350.

Sacerdote, Bruce (2011), "Peer effects in education: How might they work, how big are they and how much do we know thus far?", *Handbook of the Economics of Education*, ed. by Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, vol. 3, Elsevier, pp. 249–277.

Saiz, Albert (2010), "The geographic determinants of housing supply", *The Quarterly Journal of Economics* 125, 1253–1296.

Shiller, Robert J (1991), "Arithmetic repeat sales price estimators", *Journal of Housing Economics* 1, 110–126.

Snyder, Thomas D and Sally A Dillow (2012), "Digest of Education Statistics, 2011. NCES 2012-001.", *National Center for Education Statistics*.

Social Research Center (2019), *Panel Study of Income Dynamics [dataset]*, Data retrieved from https://simba.isr.umich.edu/data/data.aspx.

The College of William and Mary and the Minnesota Population Center (2011), *School Attendance Boundary Information System (SABINS): Version 1.0 [dataset]*, Data retrieved from http://www.sabinsdata.org.

U.S. Bureau of Labor Statistics (2021), *"Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [dataset]*, Data retrieved from FRED (Federal Reserve Bank of St. Louis): https://fred.stlouisfed.org/series/CPIAUCSL.

U.S. Census Bureau (2019), *2000-2014 American Community Survey 5-year Public Use Microdata Samples [dataset]*, Data retrieved from https://data.census.gov/.

Zheng, Angela and James Graham (2021), "Public Education Inequality and Intergenerational Mobility", *American Economic Journal: Macroeconomics*.

Zillow (2020), *Zillow Transaction and Assessment Dataset [dataset]*.

Online Appendix

A. Empirical Analysis Details

This appendix provides additional details about our sample construction, estimation of school and teacher value-added (VA), and construction of our house price instrument.

A.1. Education Sample Construction

Our data cover elementary grades for a large urban school district for school years 2002-03 through 2016-17. Given the requirement for lagged test scores, we start with the entire enrollment history of students in the district in grades 3-5 for the school years 2003-04 through 2016-17. We then drop academic years 2013-14 and 2014-15 from the dataset along with third grade after 2012-13 due to missing data.[33] Our analysis sample therefore cover grades 4-5 from 2003-04 through 2012-13 and 2015-16 through 2016-17 school years and third grade from 2003-04 through 2012-13. These data cover roughly 800,000 students with 1.7 million student-year observations.

Our data also include detailed demographic information. Specifically, we have information about parental education (five education groups), economically disadvantaged status, ethnicity (seven ethnic groups), gender, limited English status, and age. Demographic coverage is near-universal for all demographic variables with the exception of parental education, which is missing for twenty-nine percent of the sample. Whenever demographic information is missing, we create a missing indicator for that variable.

We make several data restrictions to arrive at our final VA samples. To start, we exclude roughly 200,000 student-year observations that lack a valid current or lagged mathematics test score; these data then constitute our sample used to estimate school VA. To arrive at our teacher VA sample, we make two additional sample restrictions. First, we drop approximately 90,000 student-year observations that cannot be matched to a teacher. Second, we only include classes with more than seven but fewer than forty students with valid current and lagged mathematics scores, losing an additional 8,500 observations.

Table B.1 reports summary statistics. Our school district is majority-hispanic and consists of a relatively low-income student body with over two-thirds of students qualifying for free or reduced price lunch.[34] Columns (2) and (3) then show the samples used to estimate VA. The VA samples are similar to the full sample, although are somewhat positively selected with student test scores being about 0.02 standard deviations higher than the full sample.[35]

A.2. Constructing School Value-Added

Using the school VA sample, we estimate estimate school VA using the following equation:

$$y_{ist} = \phi X_{ist} + \mu_{st} + \epsilon_{ist} \,, \tag{A.1}$$

[33] Data are missing for 2013-14 and 2014-15 due to a change in the statewide testing regime that occurred in 2013-14, which resulted in no test score data that year and also eliminated the second grade test thereafter. As lagged test scores are required when computing value-added, we drop academic years 2013-14 and 2014-15 from the dataset, as well as third grade after 2012-13.

[34] Free or reduced price lunch eligibility is often used as a poverty indicator in education data sets as students are only eligible if their family income is at or below 185 percent of the poverty level.

[35] The positive selection is driven by the requirement that students have a lagged test score, as students without lagged test scores tend to be lower-performing. This moderate positive selection into the VA analysis sample is ubiquitous in the VA literature.

where y_{ist} is the mathematics score of student i in school s at time t, X_{ist} captures observed characteristics of the student (demographics, past academic performance, and family background), and μ_{st} is the school's contribution to student test scores in year t, or simply school VA. The error term ϵ_{ist} is assumed to be independently and identically distributed normal with variance σ_ϵ^2. A key requirement for school VA, μ_{st}, to be unbiased is that the control vector X_{ist} is sufficiently rich, with lagged test scores acting as the key control (Chetty et al., 2014a). We therefore follow this literature and include a rich set of controls in X_{ist}, including: (i) cubic polynomial in prior-year scores in mathematics and English interacted with grade dummies,[36] (ii) individual-level demographics, including parental education (five education groups), economically disadvantaged status, ethnicity (seven ethnic groups), gender, limited English status, and age interacted with grade dummies, and (iii) grade and year dummies. In contrast to much of the VA literature, however, we do not include school or school-grade level means of prior-year test scores or individual covariates so that we can decompose school VA into the portion coming from the school itself and the portion coming through peer effects (see Section 2.4).

The parameters of interest in equation (A.1), μ_{st}, can be estimated via the maximum likelihood estimator (often referred to as the fixed effect estimator) which is given by:

$$\mu_{st} = \frac{1}{n_{st}} \sum_{i=1}^{n_{st}} (y_{ist} - \hat{\phi} X_{ist}) \,, \tag{A.2}$$

where n_{st} is the total number of students in the VA sample at school s in year t.[37] While the estimator given by equation (A.2) is consistent, it is rarely used in practice due to finite sample considerations. Instead, the VA literature uses empirical Bayes methods to leverage additional information about the distribution of school VA to modify poor-quality estimates for some schools based on observations for other schools. We follow the lead of this well-developed literature and employ the parametric empirical Bayes estimator (see Morris, 1983), which takes the following form:

$$\delta_{st} = \mu_{st} \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2 / \, n_{st}} \,, \tag{A.3}$$

where σ_μ^2 and σ_ϵ^2 represent the variance of school value-added and idiosyncratic student shocks, respectively. These model parameters are estimated via maximum likelihood and then plugged-in to equation (A.3) to get our school VA estimates, $\hat{\delta}_{st}$.

A.3. Decomposing School Value-Added into Peer and Peer-Invariant VA

This subsection describes in greater detail our decomposition – using a methodology borrowed from Altonji et al. (2015) and Allende (2019) – of school VA into its peer and peer-invariant components (see Section 2.5). Formally, let VA_{st} denote the VA of school s in year t and let the vector $\boldsymbol{x}_i$ include characteristics that are assumed to have a potential impact on the outcomes of other students. Following Allende, 2019, we define $\boldsymbol{x}_i$ as a two-dimensional socioeconomic type $\boldsymbol{x}_i = (x_i^y, x_i^e)$, composed by the binary variables x_i^y and x_i^e that indicate whether

[36]When prior English test scores are missing, we set the English score to zero and include an indicator for missing data interacted with the cubic polynomial in prior-year mathematics scores.

[37]We follow much of the VA literature and estimate $\hat{\phi}$ in a first step where we regress $y_{ist} = \phi X_{ist} + \mu_s + \epsilon_{ist}$ to estimate $\hat{\phi}$ and then construct the fixed effects estimates using equation (A.2) in the second step. Alternatively, one could estimate the fixed effects in a single step, although results are near-identical. See Koedel et al., 2015 for a discussion of one- versus two-step estimators in the context of VA.

the student is socioeconomically disadvantaged and/or has educated parents. Specifically, we define a socioeconomically disadvantaged student as one who is eligible for free or reduced price lunch and students with educated parents as those whose parents are high school graduates.

We then characterize the peers in the school as a vector, $\boldsymbol{z}_{st}$, that includes the mean for the characteristics in $\boldsymbol{x}_i$ for school s at time t. We then decompose the peer and peer-invariants components of school VA by projecting (estimated) school VA, $\widehat{VA}_{st}$, onto the peers vector, $\boldsymbol{z}_{st}$, plus a school fixed effect:

$$\widehat{VA}_{st} = \boldsymbol{z}'_{st}\pi^z + \alpha_s + \epsilon_{st} \,. \tag{A.4}$$

The portion of school quality coming directly through peers, 'Peer VA,' is given by $\boldsymbol{z}'_{st}\hat{\pi}^z$. The portion of school quality not coming through peers, 'peer-invariant VA,' is then the portion of VA unexplained by peers and so is recovered by subtracting $\boldsymbol{z}'_{st}\hat{\pi}^z$ from $\widehat{VA}_{st}$.

A.4. Constructing Teacher Value-Added

Constructing Teacher Value-Added: The procedure to estimate teacher quality is near-identical to our school VA estimation procedure. Using the teacher VA sample, we estimate teacher VA using the following equation:

$$y_{ijt} = \phi X_{ijt} + \alpha_j + \epsilon_{ijt} \,, \tag{A.5}$$

where y_{ijt} is the mathematics score of student i assigned to teacher j at time t, X_{ijt} captures observed characteristics of the student (we use the same control vector as for school VA, although also include school-grade and classroom level means of prior-year test scores and individual covariates), and α_j is teacher j's (time-invariant) contribution to student test scores, or simply teacher VA. Once again, the error term ϵ_{ist} is assumed to be independently and identically distributed normal with variance σ^2_ϵ.

We then construct our estimate of teacher VA, μ_j, using the empirical Bayes estimator:

$$\mu_j = \alpha_j \frac{\sigma^2_\alpha}{\sigma^2_\alpha + \sigma^2_\epsilon / \sum_t n_{jt}} \,, \tag{A.6}$$

where $\alpha_j \equiv \sum_t \sum_{i=1}^{n_{jt}} (y_{ijt} - \hat{\phi} X_{ijt}) / \sum_t n_{jt}$ where n_{jt} is the size of the class taught by teacher j in year t. As before, σ^2_α and σ^2_ϵ represent the variance of teacher value-added and idiosyncratic student shocks, respectively. These model parameters are estimated via maximum likelihood and then plugged-in to equation (A.6) to get our teacher VA estimates, $\hat{\mu}_j$.

Calculating Turnover-Induced Teacher Value-Added Changes: The turnover-induced change in teacher VA is then calculated over the relevant time period by finding the VA of teachers that are entering and exiting a given school. Specifically, let n_{jt} denote the enrollment of teacher j's class in period t and let μ_j^{-s} denote teacher j's value-added excluding years where they taught at school s. (The exclusion of years where the teacher taught at school s ensures that the changes in teacher VA at school s solely come from teacher staffing changes and not from within-teacher quality changes.)

We then take all teachers who enter school s in period t from another school s' in $t-1$[38]

[38] The set s' also includes the option of not teaching. We therefore include teachers who enter school s but did not teach in the prior year as part of our identifying variation.

and find the enrollment-weighted VA, $\hat{Z}^{enter}_{st}$, of these teachers in school s:

$$\hat{Z}^{enter}_{st} = \frac{\sum_j n_{jt}\hat{\mu}_j^{-s}\mathbb{1}\{st \neq s', t-1\}}{\sum_j n_{jt}}. \tag{A.7}$$

Analogously, we take all teachers who exited school s in period $t-1$ and find the enrollment-weighted VA, $\hat{Z}^{exit}_{st}$, that these teachers would have contributed to school s in period t had they not left:

$$\hat{Z}^{exit}_{st} = \frac{\sum_j n_{j,t-1}\hat{\mu}_j^{-s}\mathbb{1}\{s't \neq st-1\}}{\sum_j n_{jt}}. \tag{A.8}$$

The change in VA at school s in year t, Z_{st}, is then given as the change in VA in school s coming from teachers that enter and exit school s in year t: $\hat{Z}_{st} = \hat{Z}^{enter}_{st} - \hat{Z}^{exit}_{st}$.

Note that equations (A.7) and (A.8) use jack-knife teacher VA estimates. These VA estimates are constructed by simply removing the jack-knife years from the calculation of teacher VA. Thereore, if we wish to remove years t and $t-1$ from the VA calculation, our jack-knife VA estimator, $\mu_j^{-\{t-1,t\}}$, would be:

$$\mu_j^{-\{t-1,t\}} = \alpha_j^{-\{t-1,t\}} \frac{\sigma^2_\alpha}{\sigma^2_\alpha + \sigma^2_\epsilon / \sum\limits_{\substack{t\neq t-1 \\ t \neq t}} n_{jt}}, \tag{A.9}$$

where $\alpha_j^{-\{t-1,t\}} \equiv \sum\limits_{\substack{t\neq t-1 \\ t \neq t}} \sum_{i=1}^{n_{jt}} (y_{ijt} - \hat{\phi}X_{ijt}) / \sum\limits_{\substack{t\neq t-1 \\ t \neq t}} n_{jt}$.

A.5. Constructing Cross-Walk between Census Tracts and School Zones

We construct the cross-walk using school attendance boundaries from 2015-16 (National Center for Education Statistics, 2018) and census tract files for 2010 from IPUMS.

We construct a mapping from census tracts to school zones as follows. Let $c_1 ... c_N$ be all the census tracts that intersect school zone z. Then x_z, the value for a sociodemographic characteristic x in school zone z, is a weighted average of $x_i, i = 1, ..., N$, the sociodemographic values for census tract i. Precisely, $x_z = \sum_{i=1}^N \omega_{z,i} x_i$. The weight, $\omega_{z,i}$ is the share of the school zone area z that intersects with census tract c_i. The cross-walk reports the population share of a given school zone that falls into each intersecting census tract.

A.6. Constructing Bartik Instrument for House Prices

We construct a Bartik-style instrument following Graham et al. (2023). Let $B_{z,t-5,t}$ denote the instrument for local house price growth between t and $t-5$. The instrument is constructed as the interaction between the local shares $\lambda_{z,c}$ of houses with a given characteristics c with the change in the aggregated marginal price of those characteristics $\Delta q_{c,t-5,t}$. The instrument is given by:

$$B_{z,t-5,t} = \sum_{d\in\mathcal{D}} \lambda_{z,d}\Delta q_{d,t-5,t} + \sum_{b\in\mathcal{B}} \lambda_{z,b}\Delta q_{b,t-5,t} + \sum_{h\in\mathcal{H}} \lambda_{z,h}\Delta q_{h,t-5,t} \tag{A.10}$$

where $d \in \mathcal{D}$, $b \in \mathcal{B}$, and $h \in \mathcal{H}$ denote distinct sets of house characteristics described in detail below, $\lambda_{z,c}$ is the share of houses in zone z with generic characteristic c, and $\Delta q_{c,t-5,t}$

is the 5-year change in the aggregate marginal price of a generic characteristic c. The local characteristic shares satisfy the adding up constraints $\Sigma_{c\in\mathcal{C}}\lambda_{z,c} = 1$ for each set of characteristics $\mathcal{C} \in \{\mathcal{D}, \mathcal{B}, \mathcal{H}\}$.

We use three sets of house characteristics that are widely reported in the ZTRAX data (Zillow, 2020). These characteristics are: the decade of construction $\mathcal{D} \equiv \{pre-1939, 1940-1949, 1950-1959, 1960-1969, 1970-1979, 1980-1989, 1990-1999, 2000-2009, 2009-2018\}$; the number of bedrooms $\mathcal{B} \equiv \{1, 2, 3, 4, 5+\}$; and number of bathrooms $\mathcal{H} \equiv \{1, 2, 3, 4+\}$.[39] We compute the local shares using ZTRAX data by tabulating characteristics from all unique properties sold between 1998 and 2019. We present the shares of physical characteristics for the average school zone in our sample in Table B.2 below.

In order to construct the aggregate marginal prices of house characteristics we estimate a hedonic pricing regression using the ZTRAX housing transactions data. The regression takes the form

$$p_{j,t} = \gamma_k + \sum_{d\in\mathcal{D}} q_{d,t}\mathbb{1}(d_j = d) + \sum_{b\in\mathcal{B}} q_{b,t}\mathbb{1}(b_j = b) + \sum_{h\in\mathcal{H}} q_{h,t}\mathbb{1}(h_j = h) + \eta_{j,t} \tag{A.11}$$

where $p_{j,t}$ is the price of property j in year t, and the dummy variables $\mathbb{1}(d_j = d)$, $\mathbb{1}(b_j = b)$, $\mathbb{1}(h_j = h)$ are equal to one for a property j with the relevant construction age, number of bedrooms, and number of bathrooms. We include county-level fixed effects γ_k to absorb average differences in the level of house prices across broad geographic areas. The time-varying coefficients $q_{d,t}$, $q_{b,t}$, and $q_{h,t}$ measure the marginal prices of house characteristics for decade built, number of bathrooms, and number of bedrooms, respectively. We compute 5-year changes in these marginal prices to construct the growth rates $\Delta q_{c,t-5,t}$ in Equation (A.10).

We estimate Equation (A.11) using house transactions from a broad geographic area in order to capture aggregate movements in the marginal prices of house characteristics. We use transactions for all houses in the US state in which our school district is located, but exclude all transactions from the school district itself. This is similar to the common leave-one-out estimator used for shift-share instruments, except that we exclude all sources of variation in house prices that might directly affect school zones in our district (i.e., all other zones within the district). This removes any mechanical correlation between changes in local house prices and our aggregate marginal house characteristic prices. As a result, we avoid the possibility of reverse causality between local price movements and the aggregate time-series variation in our instrument.

Let $B_{z,t-5,t}$ denote the Bartik-like instrument for local house price growth between t and $t-5$. Identification requires that the instrument $B_{z,t-5,t}$ does not affect local school quality growth except through its effects on local house price growth:

$$\text{Cov}(B_{z,t-5,t}, \epsilon_{z,t} | \alpha_z, \alpha_t, X_{z,t,t+5}) = 0$$

Following Goldsmith-Pinkham et al. (2020), we assume that that identification follows from exogeneity of the local shares embedded in our instrument. Specifically, cross-sectional variation in local housing characteristic shares $\lambda_{z,c}$ is exogenous to the error term $\epsilon_{z,t}$. In other words, unobserved shocks to local school quality must be uncorrelated with the composition of the local housing stock.

[39] Graham et al. (2023) also considers an extension of the instrument to include characteristics describing house floor size and property lot size. They find that this extended instrument provides little additional information relative to year, bedroom, and bathroom characteristics.

B. Quantitative Model Details

B.1. Model Discretization

The model statespace is given by $\mathbf{s} = \{b, y, a, n\}$. The number of grid points in each dimension are N_b, N_y, N_a, and N_n. We set the number of neighborhoods $N_n = 5$. Child ability a follows an AR(1) process as in Equation (3) with parameters μ_a, ρ_a, and σ_a. We set $N_a = 7$ and discretize the process using the Rouwenhorst method.

At age $j = 2$ adults receive log-normally distributed productivity shocks ε_y. We discretize the shocks process using a Gauss-Hermite method with $N_{\varepsilon_y} = 7$ nodes.

We set the minimum liquid asset grid size to $\underline{b} = -\theta\Delta\overline{P_n}$ where $\overline{P_n} = P_E$ is the maximum steady state house price, $\Delta = 1.10$ is the maximum gross increase in house prices observed in our dynamic experiments from Section 3.5, and θ is the maximum mortgage LTV ratio. We set the maximum liquid asset grid point equal to five times maximum possible income. We set $N_b = 50$, and we split the grid evenly between negative and positive values. Finally, we distribute grid points polynomially within the negative and positive parts of the asset space.

B.2. Scaling the Ability Process

Following the literature, we could ensure minimum housing affordability by allowing for an intensive margin of house size choice, or by assuming households can only rent and that the lowest rental rate is normalized to zero (**fogli2019end**).

One difficulty in computing equilibria of our model is that for a given income distribution, nothing guarantees that at least some houses are affordable for all households. To address this problem we normalize the mean of the child ability process μ_a to ensure that the poorest household at age $j = 1$ can afford the downpayment on a house in the least expensive neighborhood:

$$(1-\theta)\underline{P_n} \leq \underline{y_k} = \left(\underline{a_k Q_n}^{\gamma}\right)^{\frac{1}{1-\delta}} \tag{D.1}$$

where underlines denote minimum values in the model, and the right-hand side comes from the lower bound for child incomes discussed in Section B.1 above. Since we discretize the ability process using the Rouwenhorst method, the smallest value of a_k is given by the grid point:

$$\underline{a_k} = \exp\left(\log(\mu_a) - \frac{1}{2}\frac{\sigma_a^2}{(1+\rho_a)(1-\rho_a)} - \frac{\sigma_a}{\sqrt{1-\rho_a^2}}\sqrt{N_a - 1}\right) \tag{D.2}$$

Combining (D.1) and (D.2), we solve for the μ_a that ensures minimum housing affordability:

$$\mu_a = \exp\left((1-\delta)\log((1-\theta)\underline{P_n}) - \gamma\log(\underline{Q_n}) + \frac{1}{2}\frac{\sigma_a^2}{(1+\rho_a)(1-\rho_a)} + \frac{\sigma_a}{\sqrt{1-\rho_a^2}}\sqrt{N_a - 1}\right)$$

where $\underline{P_n} = P_A$ and $\underline{Q_n} = Q_A$. The parameter μ_a is then updated endogenously during the calibration process used to determine δ, γ, ρ_a, σ_a, P_A, and Q_A.

B.3. Mapping National Income Distribution to District Income Distribution

We use intergenerational mobility statistics from Chetty et al. (2018). These statistics describe the mobility of income from parents to children. For the most part, parent and child incomes are calculated in the form of income ranks. Parental income is calculated as an average of household income in the late 1990s. Child income is calculated as an average over household income in 2014 and 2015. Chetty et al. (2018) then create percentile ranks for child income and parental income in the aggregate income distribution for the U.S. for their respective cohort.

We aim to match these intergenerational mobility statistics in our model. However, since our model represents a specific school district - the 25th percentile of income in our model will not necessarily match the 25th percentile of aggregate U.S. income. To address this issue we map the percentiles from the national income distribution to percentiles from the income distribution in our school district

For child income (measured in 2014-15), we use the American Community 5-year estimates. We try to follow Chetty et al. (2018) as closely as possible: we drop negative incomes and only keep individuals aged 31 to 37. We then calculate income percentiles for the national income distribution and for the specific county that contains our school district of interest.

For parental income we use the public use microdata sample from the 2000 U.S. Census. We impute missing incomes with zeros, drop negative incomes, and only keep individuals with children. In the Chetty et al. (2018) sample, kids were born between 1978 and 1983. We assume the earliest age a parent had a child was 18, implying that the parent was 40 in 2000. We keep all households heads between 40-65.

We then convert all model ranks (school district ranks) to ranks in the national income distribution.

B.4. Moving Regression in the Data

We use the Panel Study of Income Dynamics (PSID) to estimate a regression of the likelihood of moving. We use data from 2001 to 2019 and only keep household heads who are homeowners. The PSID has a variable asking if households have moved since the prior wave (2-years ago). As moving rates vary across the country, we restrict our sample to the state that contains our school district of interest from the empirical analysis.

In our model, a time period is fifteen years. In the PSID we study the likelihood of households moving with a ten year period.[40]

We run the following regression:

$$\text{move}_{i,t} = \zeta_0 + \zeta_1 \log \text{house}_{i,t} + \log \text{income}_{i,t} + \text{age}_i + \lambda_t + u_{i,t} \tag{D.3}$$

where move_i is an indicator for if the household moved at least once over any ten year window within 2001-2019. $\log \text{house}$ is the logarithm of the house value during the start of the ten year period and $\log \text{income}$ is the median household income during the ten year period. Fixed effects for household age and year are included.

[40]Restricting to only moves within a fifteen year period would cut the sample size too short.

	(1)
	$move_{i,t}$
log house	-0.0900*
	(-2.15)
Income	0.186***
	(3.35)
_cons	-0.676
	(-1.46)
N	194
R^2	0.195

Notes: This table presents estimates of Equation (D.3) The dependent variable is $move_{i,t}$.t statistics in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.5. Parameter Identification

Following a similar analysis in Gregory et al. (2022), we investigate how the targeted statistics in Table 5 are affected by changes in the model parameters. Denote a vector of moments m, where these moments in the data are a function of a generic parameter vector Θ, and we denote $\bar{\Theta}$ as the calibrated parameters reported in Panel (b) of Table 4. Let $\mathbb{1}_i$ be a selection vector with all zeros except for a value of one in the i^{th} position. Then for each parameter Θ_i in the vector we re-solve the model under parameterizations $\Theta e^{\mathbb{1}_i \epsilon}$ and $\Theta e^{-\mathbb{1}_i \epsilon}$. We then calculate the symmetric finite difference in the moment vector:

$$\frac{m(\Theta e^{\mathbb{1}_i \epsilon}) - m(\Theta e^{-\mathbb{1}_i \epsilon})}{\frac{m(\Theta e^{-\mathbb{1}_i \epsilon})}{2\epsilon}}$$

We set $\epsilon = 0.05$, or 5 percent of the benchmark parameter values. We can then interpret the finite differences as elasticities of the moments with respect to model parameters.

We first report the sensitivity of the model with respect to all internally calibrated parameters $\{\varphi, \eta, \sigma_n, \sigma_a, \rho_a, \sigma_y, \gamma, \delta, Z_B, Z_C, Z_D, Z_E\}$. Figures 6 and 7 list the same 12 moments along the y-axis, and reports the elasticity of the moments along the x-axis. Then, in Figure 8 we also show model sensitivity to several key externally calibrated parameters: $\{\beta, \alpha, \kappa, \theta\}$.

Figure 6: Sensitivity of model statistics to internally calibrated parameters

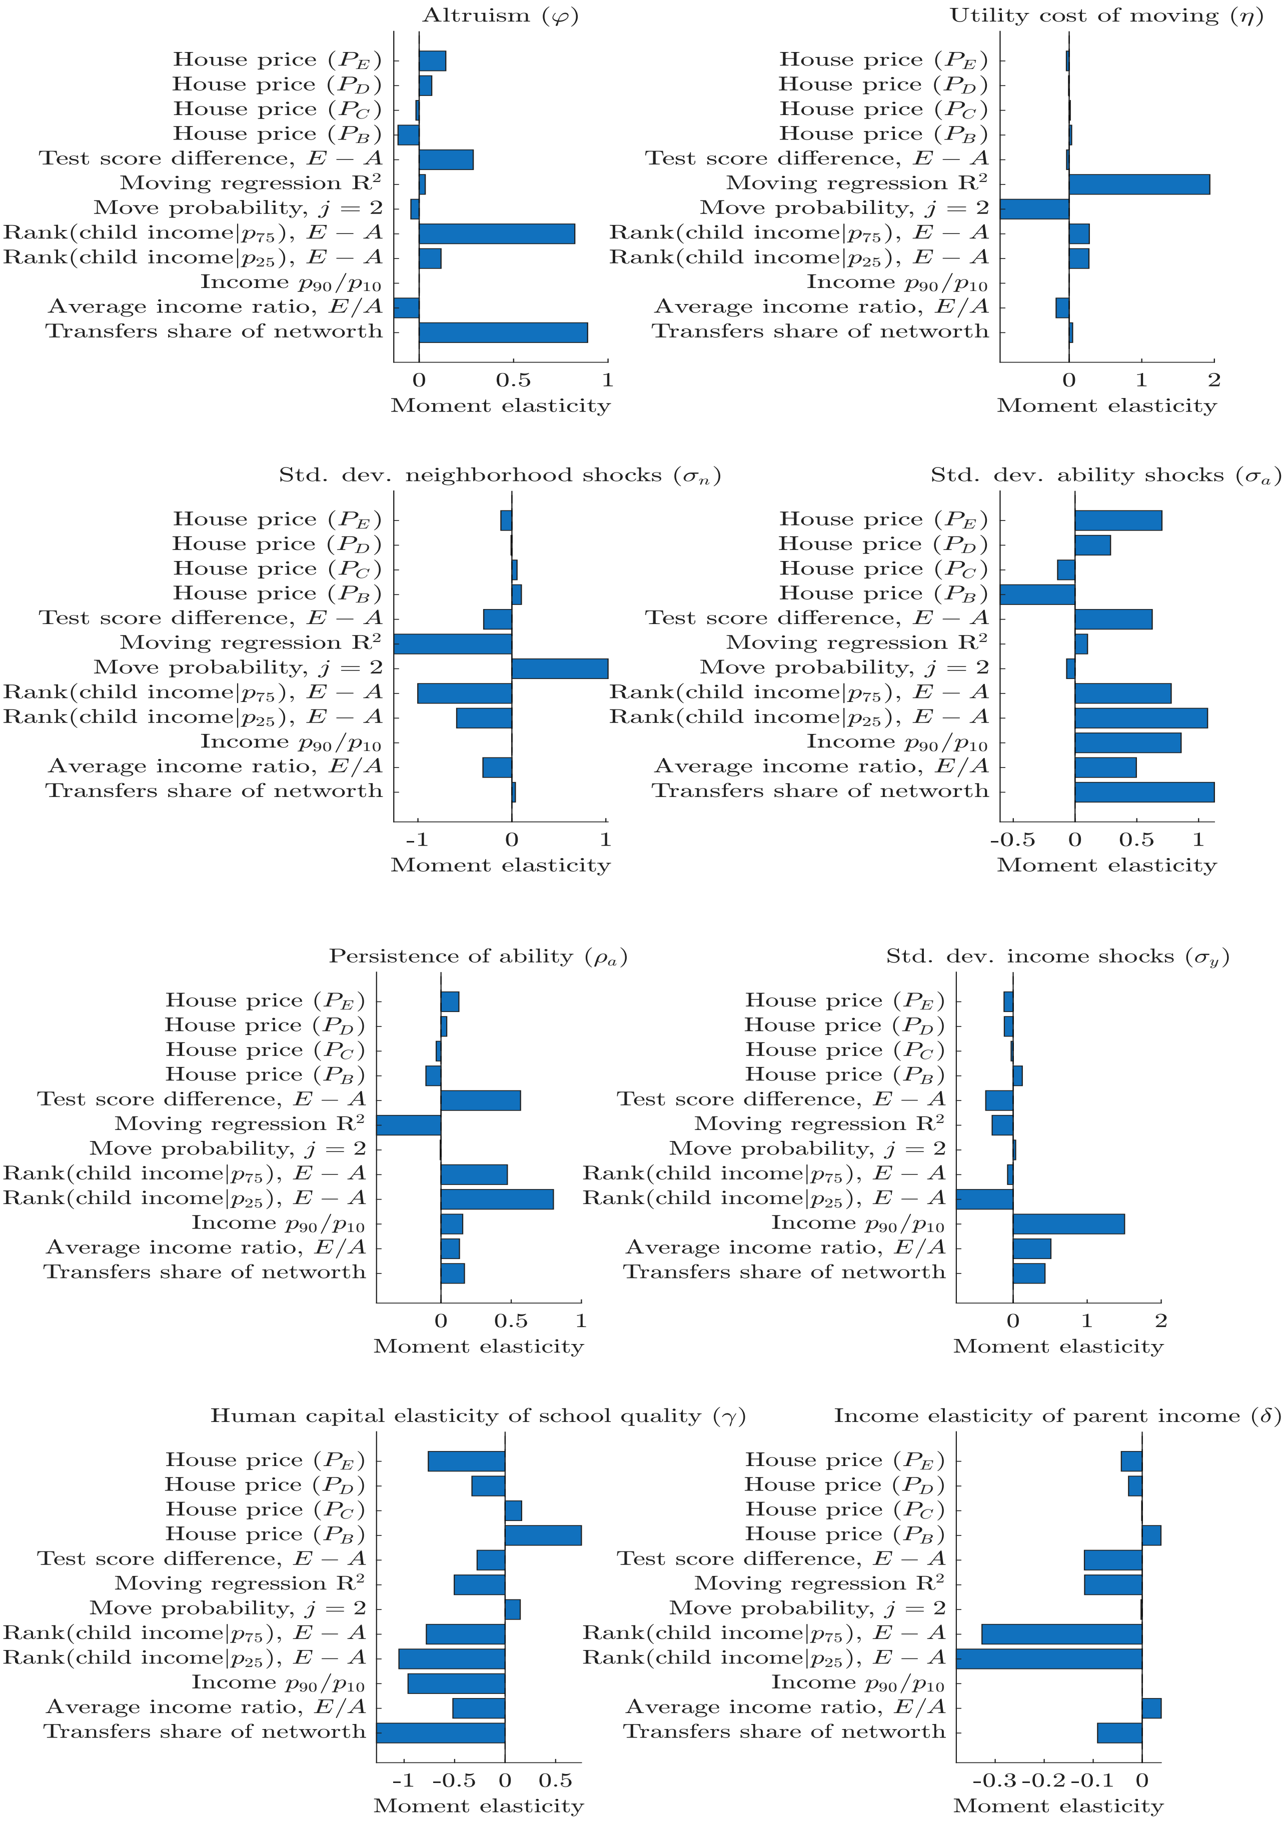

Figure 7: Sensitivity of model statistics to internally calibrated parameters

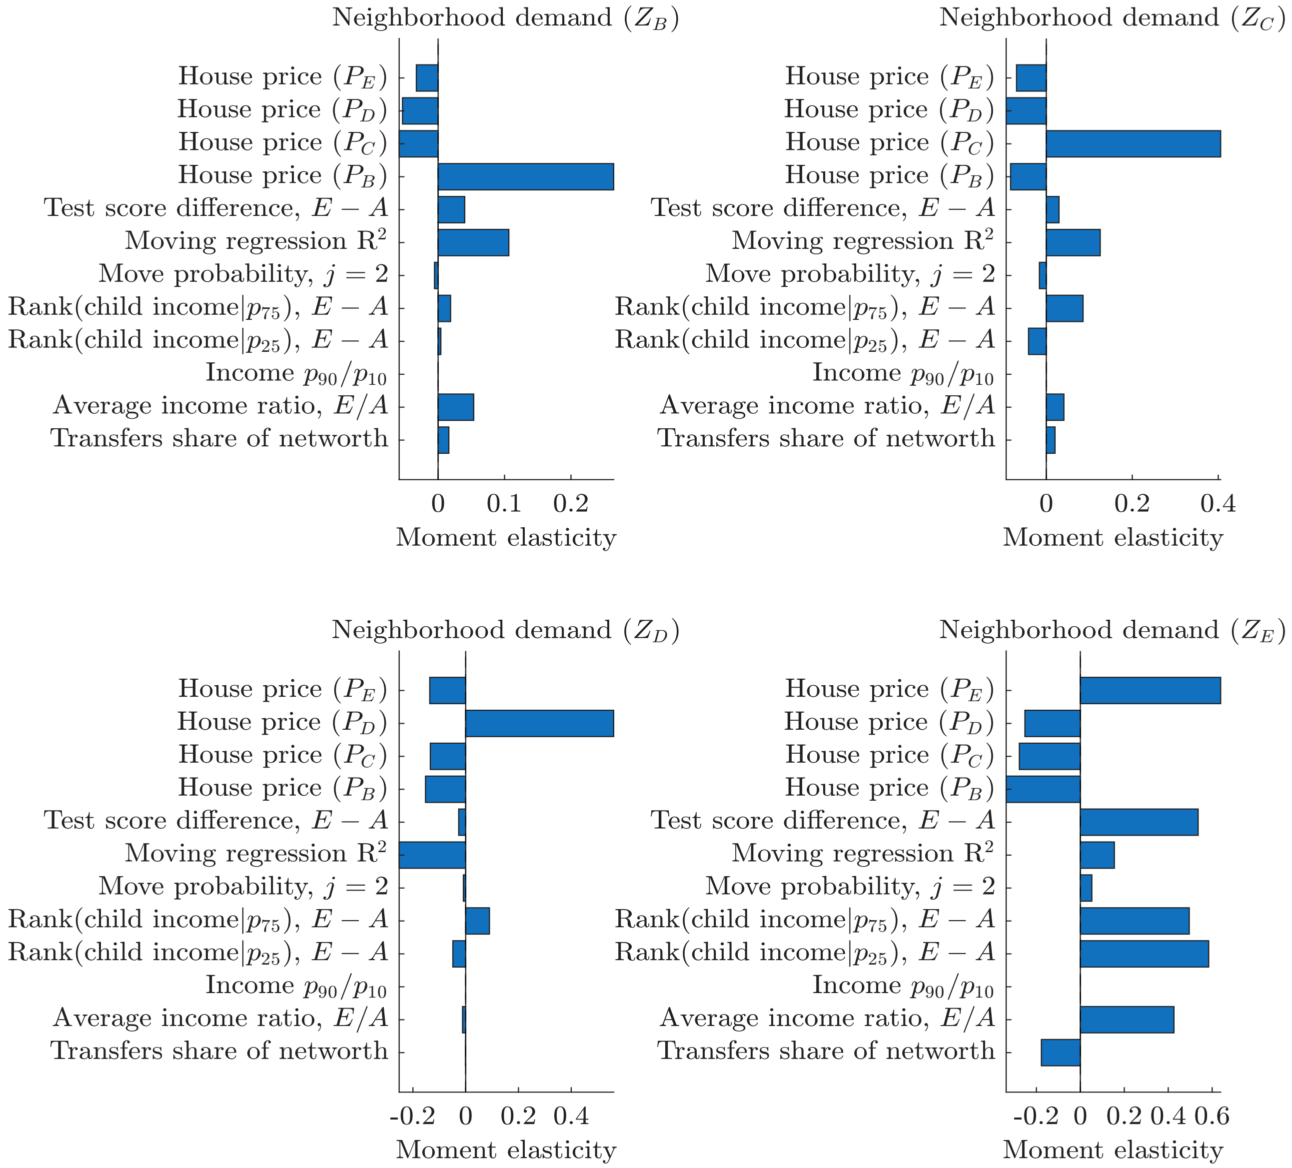

Figure 8: Sensitivity of model statistics to externally calibrated parameters

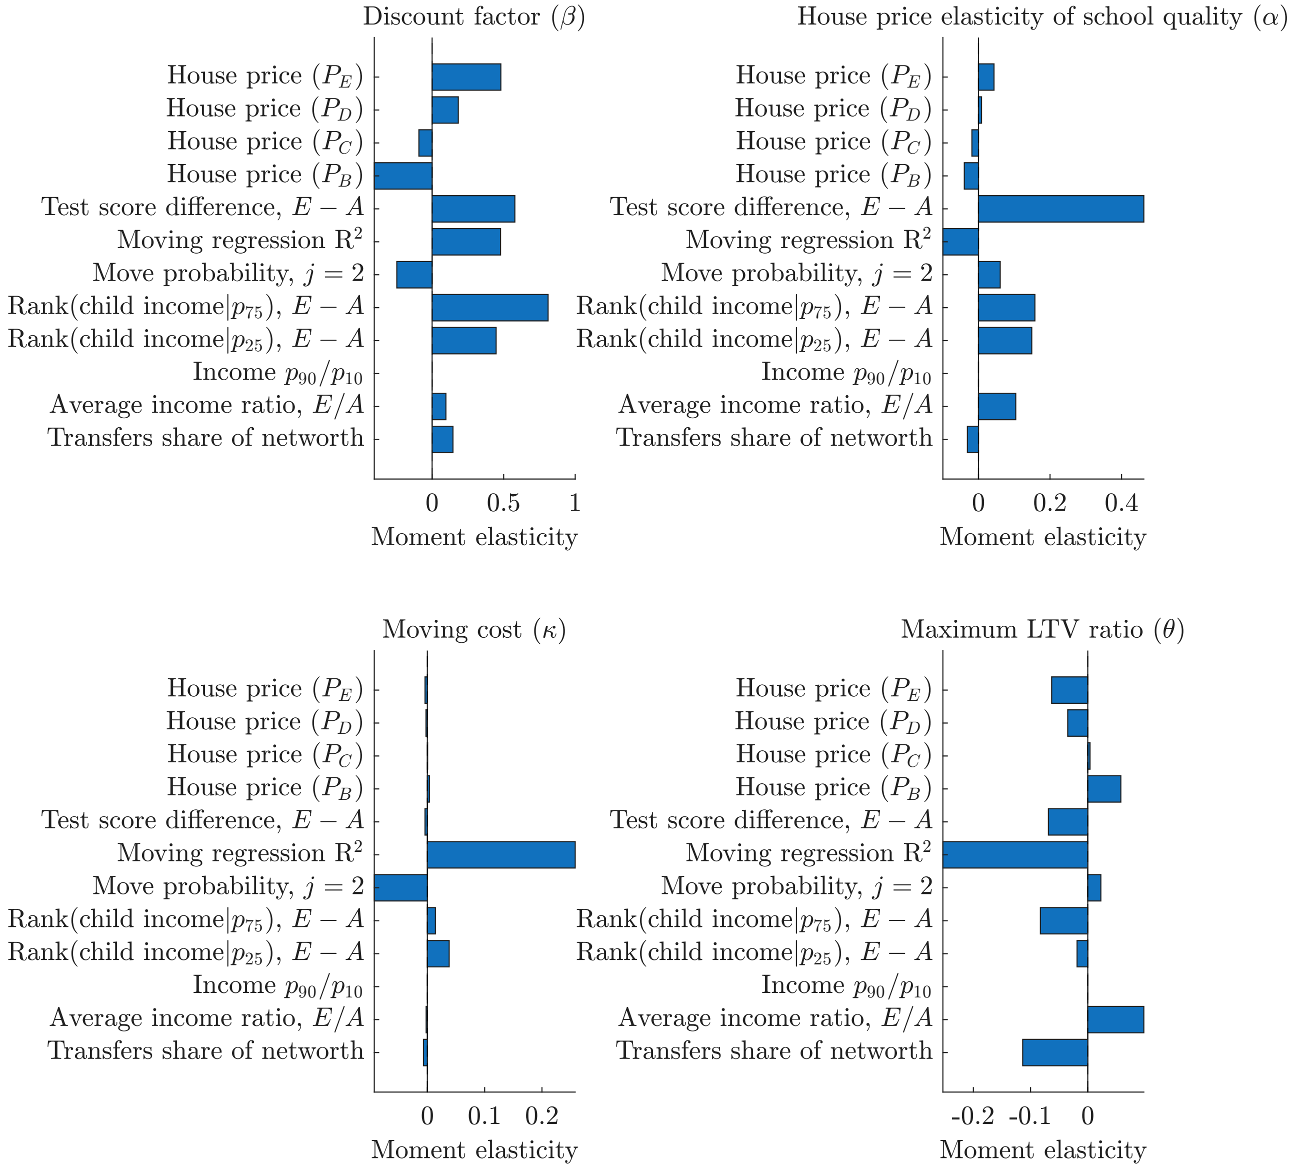

B.6. Additional Details of Dynamic Model Experiments

Figure 9: Transition Paths Following Neighborhood Demand Shock to A

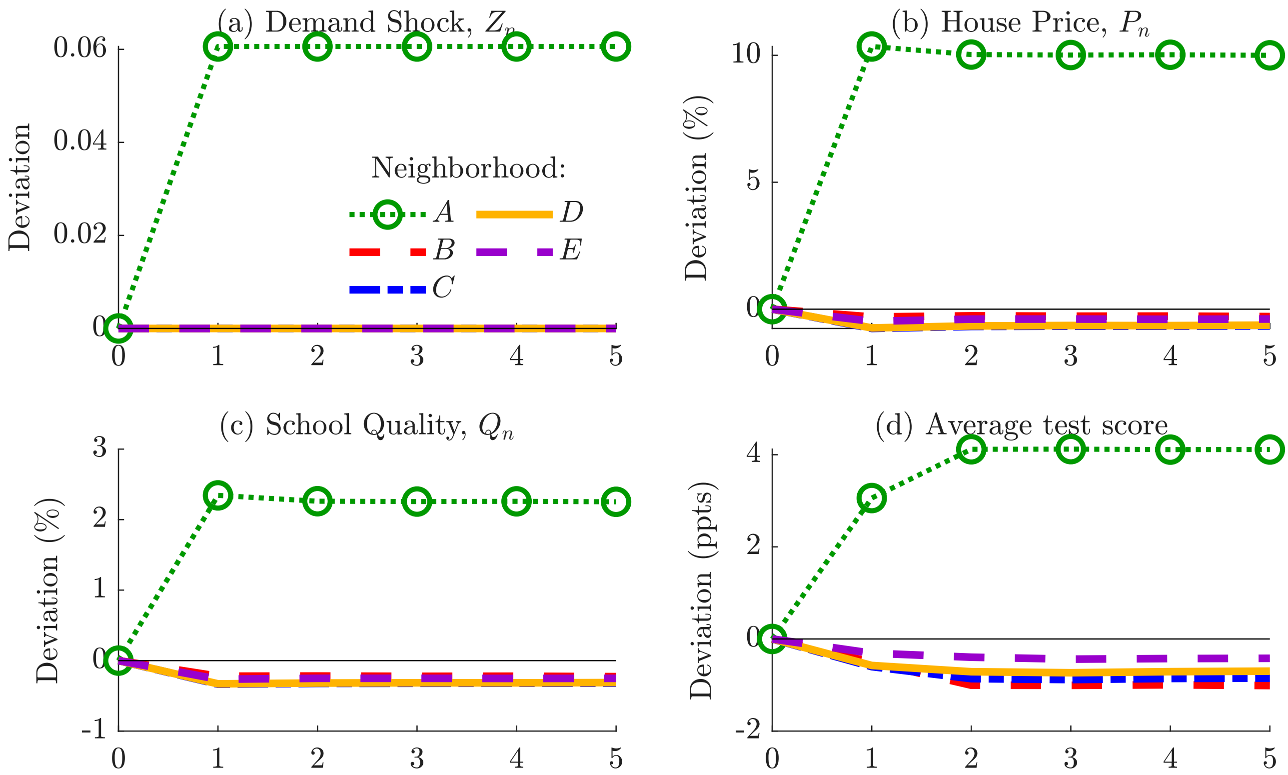

Figure 10: Transition Paths Following Neighborhood Demand Shock to B

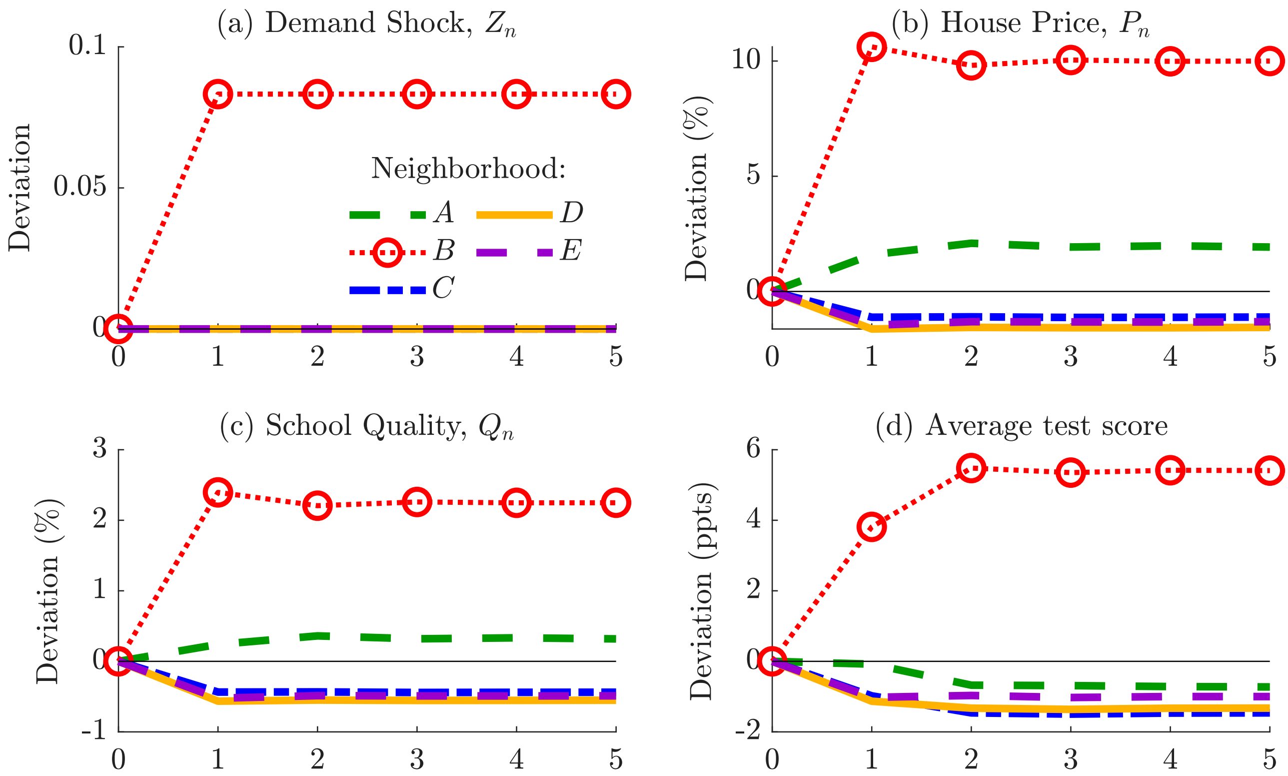

Figure 11: Transition Paths Following Neighborhood Demand Shock to C

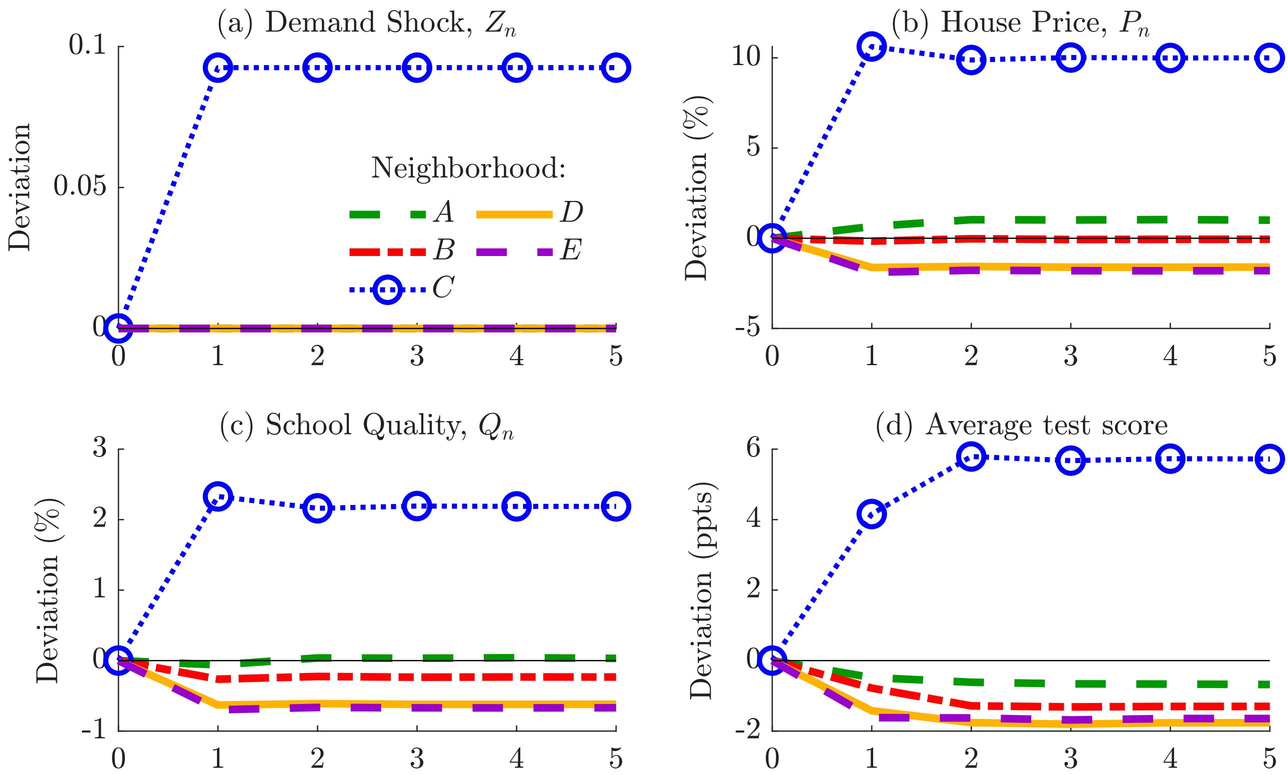

Figure 12: Transition Paths Following Neighborhood Demand Shock to D

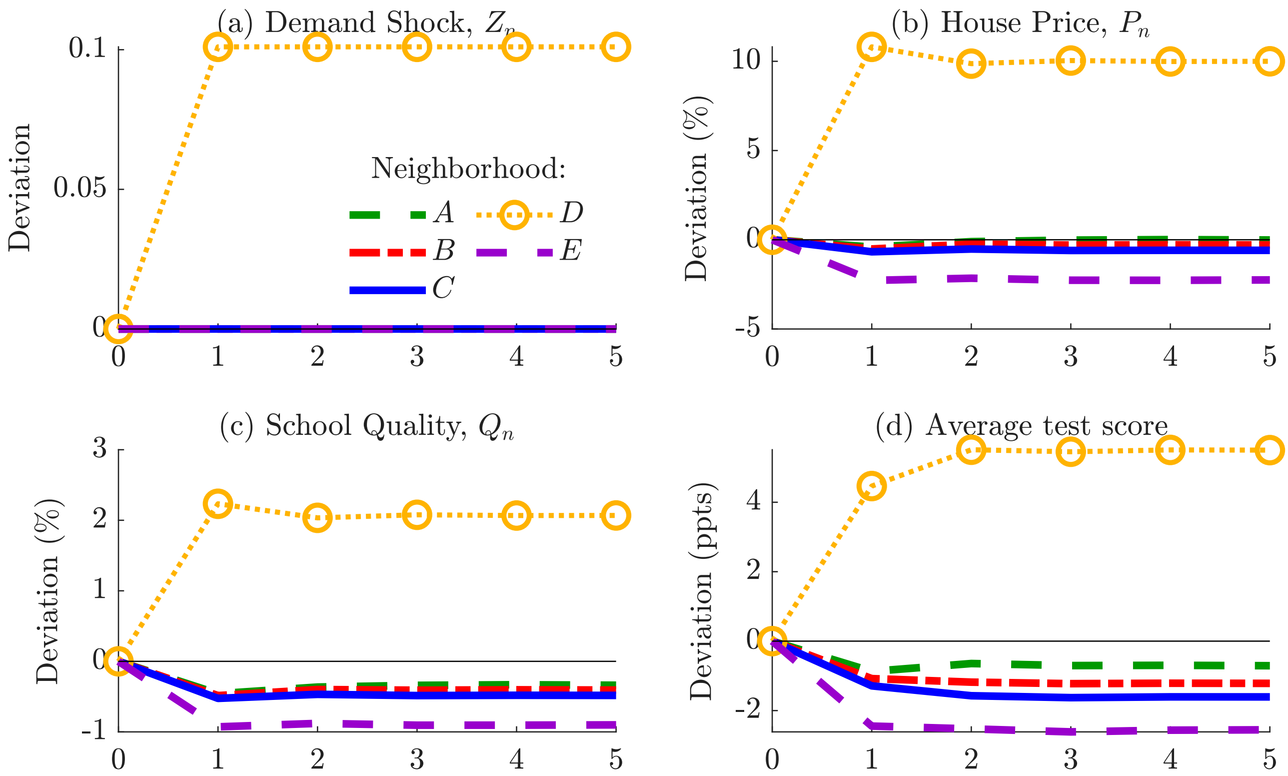

Figure 13: Transition Paths Following Neighborhood Demand Shock to E

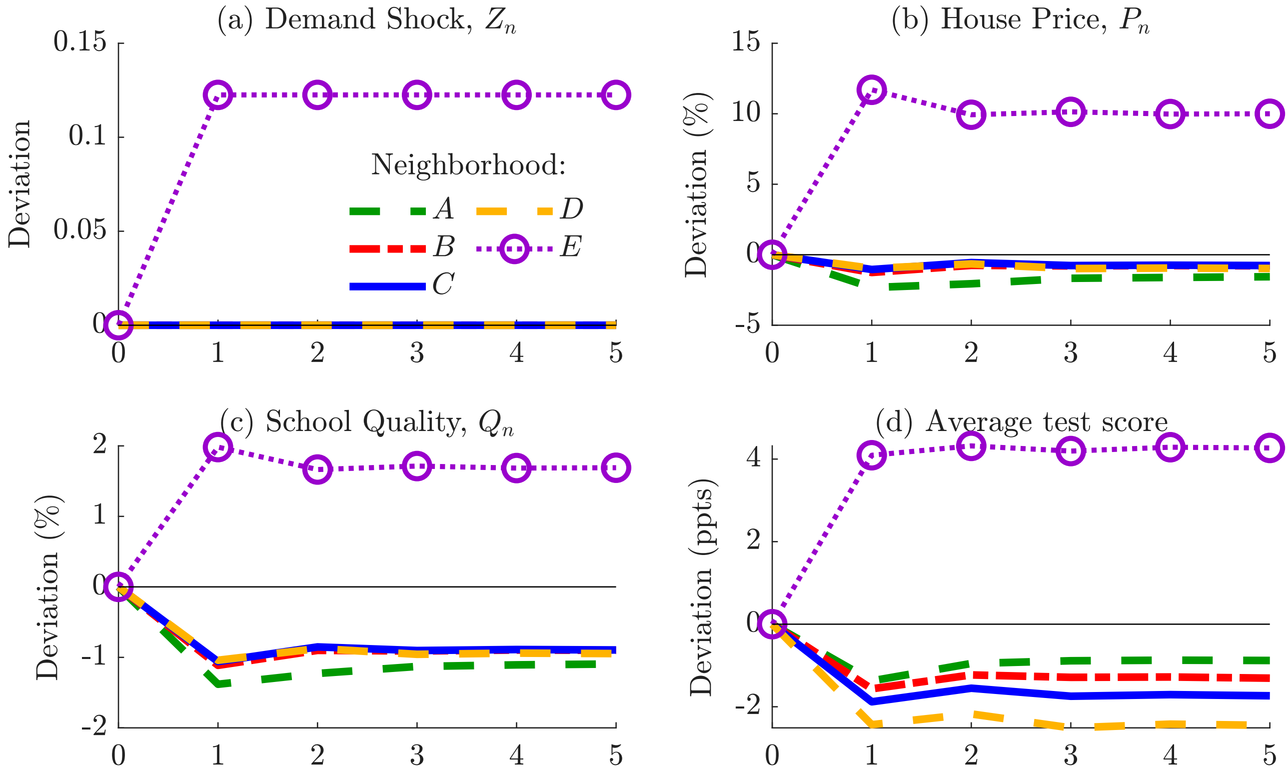

B.7. Additional Housing Market Shock Experiments

Table B.1: Housing Wealth Effects of Neighborhood Housing Supply Shocks

	Marginal propensity to consume	Marginal propensity to transfer	Marginal change in child income	Marginal propensity to move
Panel (a): All households				
All	0.023	0.011	0.018	0.011
Panel (b): Heterogeneous responses by demographics				
Low child ability	0.018	0.017	0.011	0.015
Low parent income	0.031	0.003	0.017	0.000
Low parent networth	0.024	0.007	0.028	0.028
Panel (c): Heterogeneous responses by neighborhood				
Low initial price	0.024	0.007	0.028	0.028
Price decrease	0.037	0.015	0.009	-0.008

Notes: Wealth effects estimated for age $j = 2$ households from Equation (13) and reported in annualized terms. Heterogeneous effects estimated by modifying Equation (13) to include interactions terms with price growth $\Delta P_{n,t}$.

C. Additional Tables and Figures

Table B.1: Summary Statistics for Calculating Value-Added

	Full Sample (1)	School Value-Added Sample[1] (2)	Teacher Value-Added Sample[2] (3)
Mean of Student Characteristics			
Mathematics Score (σ)	0.00	0.02	0.02
Reading Score (σ)	0.00	0.02	0.02
Lagged Mathematics Score (σ)	0.02	0.03	0.04
Lagged Reading Score (σ)	0.02	0.03	0.03
% White	9.2	9.3	8.9
% Black	9.9	9.1	9.0
% Hispanic	74.2	75.0	75.4
% Asian	4.2	4.3	4.3
% Free or Reduced Price Lunch	69.5	70.0	70.9
% English Learners	30.2	30.4	30.5
Parental Education:[3]			
% High School Dropout	34.4	34.6	34.7
% High School Graduate	45.5	45.3	45.6
% College Graduate	20.1	20.1	19.7
# of Students	839,248	743,727	717,023
# of Teachers	-	-	14,536
Observations (student-year)	1,772,731	1,558,687	1,461,842

Notes: This table presents summary statistics for the variables in our administrative education data set that we use to calculate value-added. We then compare the full sample of students in our data to the samples used to calculate school and teacher value-added.

[1] Same as the full sample, but dropping students with missing current or lagged mathematics scores.

[2] Same as the school value-added sample in column (2), but dropping students who cannot be uniquely matched to a teacher.

[3] The 'High School Graduate' category also includes parents with 'Some College,' while 'College Graduate' also incorporates those with graduate school degrees. Roughly thirty percent of observations are missing parental education data or have parental education recorded as "Decline to Answer."

Table B.2: Housing and School Zone Characteristics

Panel A: Housing Characteristics

Number of Houses	Average Sale Price	Average Bedrooms	Average # Bathrooms	Average Year Built	Median Lot Size (sq feet)	Average log House Price Change (5-yr)
717,528	386,938	2.9	2.2	1958	7500	0.23

Panel B: School Zone Demographics

% Bachelor's	Median Age	% Homeownership	% Married with Kids	% Unemployed	% Manufacturing	% Service
30	34	40	32	10	11	21

Panel C: Average School Zone Physical Characteristics Share

% Pre 1939	% 1940-1970	% 1970-2000	% Post 2000	% 1 Bedroom	% 2 Bedroom	% 3 + Bedroom
39	39	18	6.5	4.5	34	62

Notes: Panel A presents summary statistics for the sample of houses that sold in our district from 1999 to 2019. House characteristic data is from (Zillow, 2020). Panel B presents average demographics across school zones in the dataset. "% Bachelor's" refers to people with a Bachelor's degree or higher. "% Manufacturing" refers to the percentage of people who work in the manufacturing industry while "% Service" refer to the percentage of people that have an occupation in the service sector. Demographics are from the American Community Survey. Panel C presents the average percent of houses in school zones with certain characteristics that are used to construct the instrument. The first four columns refer to the time period of construction.

Table B.3: Correlation in Share of Housing Characteristics between Sales in 1999-2004 and 2014-2019

Panel (a): Bedrooms and Bathrooms	1 Bedroom	2 Bedroom	3 Bedroom	4+ Bedroom	1 Bath	2 Bath	3+ Bath
Correlation	0.69	0.87	0.80	0.88	0.93	0.82	0.94
Transactions Share	0.05	0.31	0.4	0.24	0.24	0.44	0.32
Panel (b): Decade Built	<1940	1940-1949	1950-1959	1960-1969	1970-1979	1980-1989	1990-1999
Correlation	0.96	0.95	0.98	0.98	0.93	0.93	0.82
Transactions Share	0.24	0.15	0.22	0.11	0.13	0.10	0.04

Notes: This table presents correlations between school zone-level shares of house characteristics computed for houses sold in 1998-99 through 2003-04 and houses sold in 2013-14 through 2018-19. For decade built, we focus on houses built before 2000, which consist of around 95% of transactions.
Source: Author's calculations using ZTRAX (Zillow, 2020).

Table B.4: Different Time Windows for House Price and School Value-Added Changes

	Dependent variable:				
	3-year	4-year	5-year	6-year	7-year
	(1)	(2)	(3)	(4)	(5)
Δ House Price	0.138**	0.244***	0.253***	0.317***	0.389**
	(0.066)	(0.086)	(0.062)	(0.087)	(0.179)
School Zones	395	396	393	393	390
Specification	2SLS	2SLS	2SLS	2SLS	2SLS
First-Stage F Stat	80.4	57.5	112.5	71.6	23.4
School Zone Controls	Yes	Yes	Yes	Yes	Yes
School Zone F.E.	Yes	Yes	Yes	Yes	Yes
Year F.E.	Yes	Yes	Yes	Yes	Yes
Observations	1,531	1,736	1,873	1,673	1,353

Notes: This table presents estimates of Δ log House Price from Equation (1) using different time periods of house price and school value added changes. Column (1) uses 3-year windows and Column (2) uses 4-year. In Column (3) we present our baseline estimate using a 5-year time period. Column (4) uses 6-years and Column (5) uses 7-years. All estimates are computed via 2SLS using the shift-share and BSH instruments. *p<0.1; **p<0.05; ***p<0.01

Table B.5: Teacher Turnover Across 1,3, and 5-year Horizons

	Time Horizon		
% of Teachers that:	1-year	3-year	5-year
Stay in the Same School	83.5	66.0	54.3
Leave To Another School	6.0	11.2	13.9
Leave District	10.5	22.8	31.8

Notes: The numbers in each column sum to one-hundred percent. We consider a teacher to have left the district if we do not observe them in our data after the relevant time horizon and the year after. Similarly, we consider a teacher to have switched schools if they appear in a different school after the relevant time horizon or one year later but were missing in the data after the relevant time horizon. Adding the extra year is done to account for 1-year teacher leaves (e.g., maternity leave) where the teacher leaves the data for one year, but has not truly left the school. We exclude the appropriate number of years at the end of our data period so that these 1-year leaves are consistently allowed.

Table B.6

	Dependent variable:			
	School Value-Added			
	(1)	(2)	(3)	(4)
log Relative Income	0.189***	0.189***	0.211***	0.211***
	(0.011)	(0.011)	(0.009)	(0.009)
Time Range	2010-2014	2010-2014	All Years	All Years
Year F.E.	No	Yes	No	Yes
Observations	1,880	1,880	3,298	3,298
Adjusted R^2	0.139	0.139	0.165	0.167
Note:				*p<0.1; **p<0.05; ***p<0.01

Figure 14: Relationship between House Price Growth and School Quality Growth

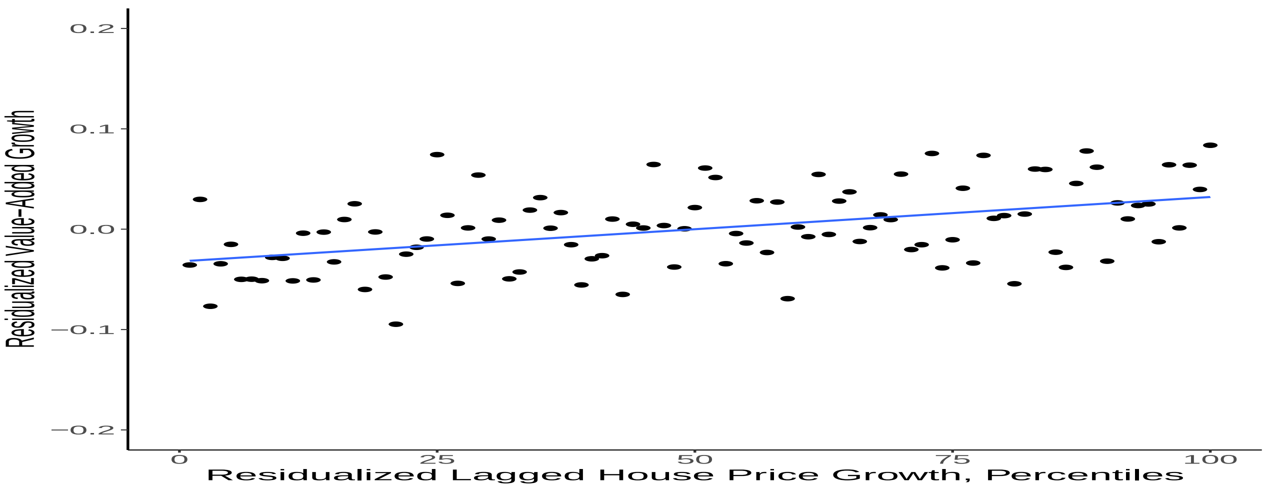

Notes: This figure plots percentiles of $\Delta \log HousePrices_{z,t-5,t}$ against $\Delta VA_{z,t,t+5}$. Both variables are residualized against school and year fixed effects. The residualized house price growths are then sorted into percentiles, and we report average growth in school VA for each bin.