TIEBOUT SORTING AND NEIGHBORHOOD STRATIFICATION [*]

Patrick Bayer
Department of Economics
Duke University and NBER

Robert McMillan
Department of Economics
University of Toronto and NBER

March 15, 2010

Abstract

Tiebout's classic 1956 paper has strong implications regarding stratification across and within jurisdictions, predicting (in the simplest instance) a hierarchy of internally homogeneous communities, ordered by household income. In practice, urban areas tend to exhibit varying degrees of within-neighborhood mixing, likely attributable to departures from several standard Tiebout assumptions – the fact that households are influenced by more than public goods packages when deciding where to live, the heterogeneous nature of the housing stock, and the role of employment geography, given commuting costs are non-zero. To shed light on the way these factors influence observed residential mixing, this paper quantifies the separate contributions of employment geography and housing preferences in reducing neighborhood stratification. It does so using an equilibrium sorting model, estimated with rich Census micro-data. Simulations based on the model using credibly-identified demand estimates show that counterfactual reductions in commuting costs lead to marked increases in education segregation and, to a lesser degree, increases in income segregation, as households now find it easier to locate in neighborhoods with similar households. In contrast, turning off preferences for housing characteristics actually reduces income segregation, indicating that the non-uniform distribution of housing serves to stratify households based on ability-to-pay. Related, we show that differences in housing also help accentuate differences in the consumption of local amenities.

Key words: Tiebout Sorting, Residential Choice, Neighborhood Stratification, Local Public Goods.

JEL: I20, H41, R21

---

[*] Thanks to Pat Bajari, Steve Berry, Tim Bresnahan, Dennis Epple, Fernando Ferreira, Caroline Hoxby, Tom Nechyba, Steve Rivkin, Jesse Rothstein, Kim Rueben, Holger Sieg, Douglas Staiger, Chris Taber, and Chris Timmins for helpful discussions about this research. We gratefully acknowledge financial support from the U.S. Department of Education, the National Science Foundation (grant SES-0137289), the Public Policy Institute of California, and SSHRC. The research in this paper was conducted while the authors were Special Sworn Status researchers of the U.S. Census Bureau, at the Berkeley and Triangle Census Research Data Centers: we would like to thank Bert Grider and Arnie Reznek for all their help. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau. This paper has been screened to ensure that no confidential data are revealed.

1 INTRODUCTION

Tiebout's classic 1956 paper emphasized the efficiency of decentralized public goods provision in a system of local jurisdictions. In so doing, it provided a provocative response to Samuelson's seminal 1954 article, which had voiced skepticism as to the efficacy of the private provision of public goods. To make this point clearly, Tiebout's treatment was stylized, focusing on a simple setting in which consumers 'voted with their feet,' choosing among exogenous public goods and taxation packages offered by a very large number of competing jurisdictions.

Beyond its enduring relevance to the debate about public goods provision, Tiebout's paper has served as an immensely fertile starting point for economists interested in understanding endogenous neighborhood stratification in an urban setting, motivated in part by concerns that excessive segregation may have adverse welfare consequences.[1] Tiebout's basic theory[2] has strong implications regarding stratification across and within jurisdictions. For example, in the simplest case where households are heterogeneous only in terms of their income, a hierarchy of internally homogeneous communities ordered by income is predicted.

In practice, urban areas tend to exhibit varying degrees of within-neighborhood mixing, and it is natural to trace this discrepancy to departures from the standard Tiebout assumptions.[3] The potential list of such departures is long. Most notably, households are heterogeneous in a number of possible ways, differing by education, income, race and family size, for instance; and these differences are likely to be associated with differences in tastes over residential choices. Related, such tastes are likely to depend on a range of locational attributes – certainly, more than just local public goods packages. The characteristics of houses, the quality levels of a variety of local amenities (including schools and crime), potentially the characteristics of neighbors, and geographic convenience are all likely to be relevant. The distribution of housing characteristics is also likely to vary markedly across a metropolitan area, given the complexity of local zoning and the durability of the housing stock, serving to sort households with different characteristics across urban space. And the geographic distribution of employment in combination with non-zero commuting costs and firm-level skill complementarities serves to spread households out across the urban area.

[1] Benabou (1993) and Cutler and Glaeser (1997) provide well-known developments of this theme.

[2] The informality of Tiebout's original presentation has prompted an important literature seeking to place the 'theory' on more solid foundations. See, for instance, Bewley (1981).

[3] The development of richer theoretical models of systems of jurisdictions has been a central preoccupation in local public finance over the past three decades. Notable here is the work of Epple and coauthors (see Epple and Zelenitz (1981), Epple, Filimon and Romer (1984, 1993), Epple and Platt (1998), and Epple and

To shed light on the way these factors influence observed residential mixing, this paper quantifies the separate contributions of employment geography and housing preferences in reducing neighborhood stratification, using a counterfactual simulation approach based on an equilibrium sorting model, estimated using rich Census data.[4] This approach builds on two strands of literature. First, long line of theoretical studies, including papers by Epple, Filimon and Romer (1984, 1993), Benabou (1993, 1996), Anas and Kim (1995), Anas (2002), Fernandez and Rogerson (1996, 1998), and Nechyba (1999, 2000), have developed and used models of sorting to analyze the way that interdependent individual decisions in the housing market aggregate up to determine the equilibrium structure of a metropolitan area. More recently, a second and related line of research has sought to take these models to the data. Notably, Epple and Sieg (1999) develop an estimator for the equilibrium sorting model of Epple, Filimon, and Romer, providing the first unified treatment of theory and empirics in the literature. In the same vein, Sieg *et al.* (2004) use this approach to explore the general equilibrium impacts of air quality improvements in the Los Angeles Basin.

In our framework, households choose residential locations, given the housing stock and employment locations of household heads, in order to maximize their utility. In terms of locational preferences, the model permits a considerable amount of household taste heterogeneity, with tastes being allowed to vary over a wide range of housing and neighborhood characteristics, including those that are endogenously determined through the sorting process; tastes also vary, potentially, according to a range of household characteristics. Neighborhood residential compositions are endogenous, and house prices adjust to equate demand with fixed supply. In equilibrium, no household can gain from moving and all local housing markets clear.

We estimate the sorting model using an econometric approach set out in Bayer, Ferreira, and McMillan (2007). This introduces unobserved neighborhood attributes into McFadden's (1978) discrete-choice housing demand model. Doing so makes clear an important endogeneity problem, as prices and neighborhood compositions, which are key choice characteristics, are likely to be correlated with these neighborhood unobservables. To address the potential endogeneity of these choice attributes, we develop instruments for price based on exogenous choice characteristics, here following Berry, Levinsohn and Pakes (1995). And to account for the potential endogeneity of school quality and sorting-dependent sociodemographics, we further

Romano (1996, 1998), among others), who show in a series of papers how housing and endogenous public goods determination can be incorporated into rigorous multi-community equilibrium models.

[4] See Bayer, McMillan, and Rueben (2009). That paper provides a formal development of the properties of the equilibrium model.

extend the boundary discontinuity approach introduced by Black (1999), making use of sharp changes in these variables in the vicinity of school attendance zone boundaries.

In taking the model to the data, we make use of restricted-access Census microdata on a very large and representative sample (1-in-7) of households in the Bay Area. The Census provides detailed information about a wide set of household characteristics, including education, income, age, family structure, and race. It also provides information about the chosen housing unit – its size, when build, whether owned, and more. Because the restricted-access version of the Census we are using specifies residential locations down to the Census block, we can merge in a great deal of additional data, relating to local amenities, land use and the characteristics of immediate neighbors. We can also construct a boundary sub-sample, based on distance to the closest school attendance zone boundary, in order to implement the boundary identification approach mentioned above. Particularly valuable for our purposes is the fact that the restricted-access version of the Census also includes information on places of work down to the block. This affords a very detailed picture of the employment geography of the metropolitan area, which we use to anchor household location decisions, taking place of work as given.[5]

Estimates of the model provide a rich characterization of heterogeneous household preferences for a variety of housing and neighborhood characteristics, including local public goods such as school quality and crime, as well as the characteristics of neighbors. Applying the identification strategies just referred to yields reasonable estimates of household willingness-to-pay for these choice attributes. As one might expect, households with higher education levels and incomes are willing to pay more for better schools, as are households with children, though the extent of the heterogeneity in willingness to pay is rather muted. While some choice characteristics are likely to be ranked similarly by all households, the rankings of others might be expected to vary depending on a household's own characteristics (their race, for example). In line with this, we find evidence of strong racial interactions in the utility function, with households of the same race showing a strong willingness to pay to live with like neighbors.

In combination with these rich preference estimates, our equilibrium model serves as a useful device for exploring the implications of changes in model primitives for residential stratification and household consumption levels. In that vein, the main part of the analysis in this paper consists of a series of counterfactual simulations intended to shed light on the factors that lead, in practice, to residential mixing.

[5] It is feasible, if more challenging, to further endogenize the place of work in the household choice process, though we defer that extension to future work.

The first set of simulations show that counterfactual reductions in commuting costs lead to marked increases in education segregation and, to a lesser degree, increases in income segregation, as households now find it easier to form neighborhoods with like households. In contrast, turning off preferences for housing characteristics actually reduces income segregation, while education segregation increases, indicating that the non-uniform distribution of housing characteristics serves to stratify households based on ability to pay. Further, we show that differences in housing help accentuate differences in the consumption of local amenities.

The rest of the paper is organized as follows: the next section sets out the equilibrium sorting model and its main properties, drawing on the treatment in Bayer *et al.* (2009). In Section 3, we discuss the estimation of the model. (Here, we borrow from Bayer, Ferreira and McMillan (2007).) Section 4 describes our detailed Census microdata from the Bay Area, and Section 5, the model estimates. Section 6 summarizes our simulation approach, before presenting the simulation findings, and Section 7 concludes.

2 THEORETICAL MODEL

In this paper, we use a version of the equilibrium model of residential sorting developed in Bayer, McMillan, and Rueben (2009). The model that we use here consists of two key elements: the individual household residential location decision problem and a market-clearing condition. While maintaining a simple structure, the model allows households to have heterogeneous preferences defined over housing and neighborhood attributes in a flexible way; it also allows housing prices and neighborhood sociodemographic compositions to be determined in equilibrium. In estimating the model, we are careful to account for the correlation that arises between unobserved housing and neighborhood attributes and both housing prices and neighborhood composition as a result of sorting. Having estimated the model, we then use it to conduct counterfactual simulations designed to measure the extent to which various housing, neighborhood, and geographic factors contribute to residential stratification on the basis of income and education.

The Residential Location Decision. We model the residential location decision of each household as a discrete choice of a single residence from a set of house types available in the market. Let X_h represent the observable characteristics of housing choice h, including characteristics of the house itself (e.g., size, age, and type), its tenure status (rented vs. owned), and the characteristics of its neighborhood (e.g., school quality, crime, land use, and topography). We use the notation capital letter Z_h to represent the average sociodemographic characteristics of the corresponding neighborhood, writing it separately from the other housing and neighborhood attributes to make explicit the fact that these characteristics are determined in equilibrium. Let p_h denote the price of housing choice h and, finally, let d_h^i denote the distance from residence h to the primary work location of household i. Each household chooses its residence h to maximize its indirect utility function V_h^i:

$$(1) \qquad \underset{(h)}{Max} \quad V_h^i = \alpha_X^i X_h + \alpha_Z^i Z_h - \alpha_p^i p_h - \alpha_d^i d_h^i + \xi_h + \varepsilon_h^i.$$

The error structure of the indirect utility is divided into a correlated component associated with each housing choice that is valued the same by all households, ξ_h, and an individual-specific term, $\varepsilon^i{}_h$. A useful interpretation of ξ_h is that it captures the unobserved quality of each housing choice, including any unobserved quality associated with its neighborhood.

Each household's valuation of choice characteristics is allowed to vary with its own characteristics, z^i, including education, income, race, employment status, and household composition. Specifically, each parameter associated with housing and neighborhood characteristics and price, α^i_j, for $j \in \{X, Z, d, p\}$, varies with a household's own characteristics according to:

$$(2) \qquad \alpha_j^i = \alpha_{0j} + \sum_{k=1}^{K} \alpha_{rj} z_k^i \ ,$$

with equation (2) describing household i's preference for choice characteristic j.

This specification of the utility function gives rise to a horizontal model of sorting in which household preferences are defined distinctly over each choice characteristic, including both school quality and neighborhood sociodemographic characteristics. This contrasts with vertical models, which restrict households to have preferences over a single locational index, thereby constraining households to have the same preference ordering across locations. The additional flexibility of the horizontal model is especially relevant for this paper as it is the magnitude of the heterogeneity in preferences for various locational factors that will determine the extent of stratification across neighborhoods.

Characterizing the Housing Market. We assume that the housing market can be fully characterized by a set of housing types that is a subset of the full set of available houses, letting the supply of housing of type h be given by S_h. We also assume that each household observed in the sample represents a continuum of households with the same observable characteristics, with the distribution of idiosyncratic tastes ε^i_h mapping into a set of choice probabilities that characterize the distribution of housing choices that would result for the continuum of households with a given set of observed characteristics.

Given the household's problem described in equations (1)-(2), household i chooses housing type h if the utility that it receives from this choice exceeds the utility that it receives from all other possible house choices - that is, when

$$(3) \qquad V^i_h > V^i_k \quad \Rightarrow \quad W^i_h + \varepsilon^i_h > W^i_k + \varepsilon^i_k \quad \Rightarrow \quad \varepsilon^i_h - \varepsilon^i_k > W^i_k - W^i_h \qquad \forall \quad k \neq h$$

where W^i_h includes all of the non-idiosyncratic components of the utility function V^i_h. As the inequalities in (3) imply, the probability that a household chooses any particular choice depends in general on the characteristics of the full set of possible house types. Thus the probability P^i_h that household i chooses housing type h can be written as a function of the full vectors of housing and neighborhood characteristics (both observed and unobserved) and prices $\{\mathbf{X}, \mathbf{Z}, \mathbf{p}, \boldsymbol{\xi}\}$:[6]

$$(4) \qquad P^i_h = f_h(z^i, \mathbf{Z}, \mathbf{X}, \mathbf{p}, \boldsymbol{\xi})$$

as well as the household's own characteristics z^i.

Aggregating the probabilities in equation (5) over all observed households yields the predicted demand for each housing type h, D_h:

$$(5) \qquad D_h = \sum_i P^i_h \, .$$

In order for the housing market to clear, the demand for houses of type h must equal the supply of such houses and so:

$$(6) \qquad D_h = S_h, \quad \forall h \quad \Rightarrow \quad \sum_i P^i_h = S_h \quad \forall h \, .$$

[6] For the purposes of characterizing the equilibrium properties of the model, we include an individual's employment location in z^i and the residential location in X_h.

Given the decentralized nature of the housing market, prices are assumed to adjust in order to clear the market. The implications of the market clearing condition defined in equation (6) for prices are very standard, with excess demand for a housing type causing price to be bid up and excess supply leading to a fall in price. Given the indirect utility function defined in (1) and a fixed set of housing and neighborhood attributes, Bayer, McMillan, and Rueben (2009) show that a unique set of prices (up to a scale) clears the market.

Given that some neighborhood attributes are endogenously determined by the sorting process itself, we define a sorting equilibrium as a set of residential location decisions and a vector of housing prices such that the housing market clears and each household makes its optimal location decision given the location decisions of all other households. In equilibrium, the vector of neighborhood sociodemographic characteristics along with the corresponding vector of market clearing prices must give rise to choice probabilities in equation (4) that aggregate back up to the same vector of neighborhood sociodemographics.[7] Whether this model gives rise to multiple equilibria depends on the distributions of preferences and available housing choices as well as the utility parameters. In general, it is not possible to establish that the equilibrium is unique *a priori*. However, estimation of the model does not require the computation of an equilibrium nor uniqueness more generally, as we describe in the next section.

3 ESTIMATION

Estimation of the model follows a two-stage procedure closely related to that developed in Berry, Levinsohn, and Pakes (1995). This section outlines the estimation procedure; a rigorous presentation is contained in Bayer, Ferreira, and McMillan (2007). It is helpful in describing the estimation approach to first introduce some notation. In particular, rewrite the indirect utility function as:

$$(7) \qquad V_h^i = \delta_h + \lambda_h^i + \varepsilon_h^i$$

where

$$(8) \qquad \delta_h = \alpha_{0X} X_h + \alpha_{0Z} Z_h - \alpha_{0p} p_h + \theta_{bh} + \xi_h$$

and

[7] Bayer, McMillan, and Rueben (2009) establish the existence of a sorting equilibrium as long as (i) the indirect utility function shown in equation (2) is decreasing in housing prices for all households; (ii) indirect utility is a continuous function of neighborhood sociodemographic characteristics; and (iii) $\boldsymbol{\varepsilon}$ is drawn from a continuous density function.

$$(9) \qquad \lambda_h^i = \left(\sum_{k=1}^{K} \alpha_{kX} z_k^i\right) X_h + \left(\sum_{k=1}^{K} \alpha_{kZ} z_k^i\right) Z_h - \left(\sum_{k=1}^{K} \alpha_{kp} z_k^i\right) p_h - \left(\sum_{k=1}^{K} \alpha_{kd} z_k^i\right) d_h .$$

In equation (78, δ_h captures the portion of utility provided by housing type h that is common to all households, and in (9), k indexes household characteristics. When the household characteristics included in the model are constructed to have mean zero, δ_h is the mean indirect utility provided by housing choice h. The unobservable component of δ_h, ξ_h, captures the portion of unobserved preferences for housing choice h that is correlated across households, while $\varepsilon_h{}^i$ represents unobserved preferences over and above this shared component.

The first step of the estimation procedure is equivalent to a Maximum Likelihood estimator applied to the individual location decisions, taking prices and neighborhood sociodemographic compositions as given. This returns estimates of the heterogeneous parameters in λ and mean indirect utilities, δ_h. The estimator is based simply on maximizing the probability that the model correctly matches each household observed in the sample with its chosen house type. In particular, for any combination of the heterogeneous parameters in λ and mean indirect utilities, δ_h, the model predicts the probability that each household i chooses house type h. We assume that $\varepsilon_h{}^i$ is drawn from the extreme value distribution, in which case this probability can be written:

$$(10) \qquad P_h^i = \mathbf{exp}(\delta_h + \hat{\lambda}_h^i) \Big/ \sum_k \mathbf{exp}(\delta_k + \hat{\lambda}_k^i)$$

Maximizing the probability that each household makes its correct housing choice gives rise to the following quasi-log-likelihood function:

$$(11) \qquad \tilde{\ell} = \sum_i \sum_h I_h^i \, \mathbf{ln}(P_h^i)$$

where $I^i{}_h$ is an indicator variable that equals 1 if household i chooses house type h in the data and 0 otherwise. The first stage of the estimation procedure consists of searching over the parameters in λ and the vector of mean indirect utilities to maximize $\tilde{\ell}$.

The Endogeneity of School Quality and Neighborhood Sociodemographic Composition. Having estimated the vector of mean indirect utilities in the first stage of the estimation procedure, the second stage involves decomposing $\boldsymbol{\delta}$ into observable and unobservable components according to the regression equation (8).[8] In estimating equation (8), important endogeneity problems need to be confronted. To the extent that house prices partly capture house and neighborhood quality unobserved to the econometrician, so the price variable will be endogenous. Estimation via least squares will thus lead to price coefficients biased towards zero, producing misleading willingness-to-pay estimates for a whole range of choice characteristics. This issue arises in the context of any differentiated products demand estimation and we follow the approach described in Bayer *et al.* (2007) to instrument for price.

A second identification issue involves the correlation of neighborhood sociodemographic characteristics $\boldsymbol{Z}$ and school quality with unobserved housing and neighborhood quality, ξ_h. To properly estimate preferences in the face of this endogeneity problem, we adapt a technique previously developed by Black (1999). Black's strategy makes use of a sample of houses near school attendance zone boundaries, estimating a hedonic price regression that includes boundary fixed effects. Intuitively, the idea is to compare houses in the same local neighborhood but on opposite sides of the boundary, exploiting the discontinuity in the right to attend a given school. Differences in valuation will then reflect differences in school quality, controlling for other neighborhood characteristics (both observed and unobserved).

As shown in Bayer, Ferreira, and McMillan (2007), however, households clearly sort with respect to these boundaries. Thus, while the boundary fixed effects are likely to control well for differences in unobserved fixed factors, neighborhood sociodemographics are likely to vary discontinuously at the boundary. This is important: it implies that boundary fixed effects isolate variation in both school quality and neighborhood sociodemographics in a small region in which unobserved fixed features (e.g., access to the transportation network) are likely to vary only slightly, thereby providing an appealing way to account for the correlation of both school quality *and* neighborhood sociodemographics with unobservable neighborhood quality.

We incorporate school attendance zone boundary fixed effects θ_{bh} when estimating equation (8). In particular, we create a series of indicator variables for each Census block that equal one if the block is within a given distance of each unique school attendance zone boundary in the metropolitan area. Bayer, Ferreira, and McMillan (2007) provides extensive descriptive evidence regarding the significant extent that the school quality and neighborhood demographics vary at these boundaries and, importantly, the essentially continuous way that other housing and neighborhood attributes run through the boundary.

[8] Notice that the set of observed residential choices provides no information that distinguishes the components of $\boldsymbol{\delta}$. That is, however $\boldsymbol{\delta}$ is broken into components, the effect on the probabilities shown in equation (10) is identical.

4 DATA

The analysis conducted is based on restricted Census data from 1990 that provide the detailed individual, household, and housing variables found in the public-use version of the Census, but also include information about the location of individual residences and workplaces at a very disaggregate level. We use data from six contiguous counties in the San Francisco Bay Area: Alameda, Contra Costa, Marin, San Mateo, San Francisco, and Santa Clara.

The Census provides a wealth of data on the individuals in the sample – race, age, educational attainment, income from various sources, household size and structure, occupation, and employment location. In addition, it provides a variety of housing characteristics: whether the unit is owned or rented, the corresponding rent or owner-reported value, number of rooms, number of bedrooms, type of structure, and the age of the building. We use these housing characteristics directly, and also construct neighborhood variables, such as neighborhood racial, education and income distributions, based on the households within the same Census block group (a Census region containing approximately 500 housing units). We merge additional data describing local conditions with each house record, constructing variables related to crime rates, land use, local schools, topography, and urban density. The list of the principal housing and neighborhood variables used in the analysis, along with means and standard deviations, is given in the first two columns of Table 1.[9]

School Boundaries. In order to implement the boundary discontinuity design, we gathered school attendance zone maps for as many elementary schools as possible in the Bay Area, for the period around the 1990 Census. Our final attendance zone sample consists of 195 elementary schools – just under a third of the total number in the Bay Area. From this boundary sample, we excluded portions of boundaries coinciding with school district boundaries, city boundaries, or large roads, since they could potentially confound our identification strategy.

For our main boundary analysis, we focus on houses in all Census blocks that are within 0.20 miles of the closest school attendance zone boundary. The average distance to the boundary for this subsample is thus quite a lot smaller than 0.20 miles. For comparison, we also analyze a further subsample, consisting of houses assigned to Census blocks within 0.10 miles of the closest attendance zone boundary. Although the 0.10-mile subsample includes approximately

[9] For each of these measures, a detailed description of the process by which the original data were assigned to each house is provided in the Data Appendix to Bayer et al. (2007).

half the number of observations, it provides a closer approximation to the ideal comparison of houses on the opposite sides of the same street, though in separate attendance zones.

Table 1 displays descriptive statistics for various samples related to the boundaries. The first two columns report means and standard deviations for the full sample while the third column reports means for the sample of houses within 0.2 miles of a school district boundary.[10] Comparing the first column to the third column of the table, the houses near school attendance zone boundaries are reasonably representative of those in the Bay Area as a whole. The fourth and fifth columns report means for houses within 0.2 miles of a boundary, comparing houses on the high versus low average test score side of the each boundary; the seventh column reports t-tests for the difference in means. Comparing these differences reveals that houses on the high side cost $18,700 more (on a mean of $250,000) and are assigned to schools with test scores that are 74 points higher on average. Moreover, houses on the high quality side of the boundary are much more likely to be inhabited by white households and households with more education and income. These types of across-boundary differences in sociodemographic composition are what one would expect if households sort on the basis of preferences for school quality. While far less significant, other housing characteristics do vary across the boundaries as well. Consequently, we expect the use of boundary fixed effects to do a reasonably good (but not perfect) job of controlling for much of the variation in unobserved housing and neighborhood quality, thereby giving rise to more accurate estimates of preferences for neighborhood sociodemographics and school quality.

5 PREFERENCE ESTIMATES

The estimation of the full model returns over 150 parameter estimates that characterize heterogeneous household preferences for housing and neighborhood attributes as well as commuting distance. While we presented some aspects of these parameter estimates in Bayer et al (2007), the focus of that paper was on the impact of including boundary fixed effects on the estimates of preferences school quality and neighborhood demographics. A much richer set of parameters is relevant for this paper, however, as the counterfactuals that we conduct below consider the impact of employment geography and heterogeneous housing preferences on the extent of residential stratification. To this end, we present a much broader set of parameter estimates here.

[10] Table 1 is taken directly from Bayer *et al.* (2007). We provide it here to aid in the readability of this paper.

Preference Estimates

Table 2 reports estimates of mean preferences for key housing and neighborhood characteristics for two specifications of equation (8). The second column shows the impact of including boundary fixed effects on the estimates of mean preferences. Comparing these columns reveals the pattern of results that one would expect if boundary fixed effects control in part for unobserved neighborhood quality and unobserved quality is positively correlated with neighborhood income and education and negatively correlated with the fraction of non-white households. Key parameter estimates include a MWTP of just under $100 for newly constructed housing and for each additional room and just over $50 per each additional mile of commuting distance.

Table 3 reports the implied estimates of the heterogeneity in MWTP for a wide variety of housing and neighborhood attributes for our preferred specification, which includes boundary fixed effects. The estimates of the heterogeneity in the MWTP for neighborhood sociodemographic characteristics reveal that while all households prefer to live in higher-income neighborhoods, *conditional on neighborhood income* households prefer to self-segregate on the basis of both race and education. In particular, the estimates imply that college-educated households are willing to pay $58 per month more than those without a college degree to live in a neighborhood that has 10 percent more college-educated households. When combined with the estimated mean MWTP of $10 per month reported in the first row, this estimate implies that households at each level of educational attainment prefer neighbors with like education levels: while college-educated households would pay an addition $32 per month to live in a neighborhood that had 10 percent more college-educated households, households without a college degree would actually need *compensating* to live in a neighborhood with 10 percent more college-educated neighbors, to the tune of $26 per month.

Similarly, the heterogeneity estimates imply that blacks are willing to pay $98 more per month than whites to live in a neighborhood that has 10 percent more black versus white households. The mean MWTP for such an increase is -$10.5 per month, primarily reflecting the negative valuation of the white majority. Thus $98 is the difference between the *positive* MWTP of black households for this change and the *negative* MWTP of white households, indicating that households have strong self-segregating racial preferences.[11]

[11] It is also important to point out that these interactions pick up any direct preferences for living near others of the same race (e.g., a recent immigrant from China may want to interact with neighbors who also have immigrated from China) as well as any unobservable neighborhood or housing amenities valued more

In contrast to education and race, neighborhood income is a normal good for all households, perhaps in part because higher income neighborhoods are better able to maintain their properties. Not surprisingly, households with children have stronger preferences for larger and older houses as well as higher quality schools and the demand for owner-occupancy increases sharply with income, education and for white and especially Asian households. The latter results are likely due in large part to the relative wealth of white and Asian households compared to that of their black and Hispanic counterparts.

6 SIMULATIONS

We begin this section by describing our simulation procedure, before discussing the construction of the exposure rate and consumption measures we focus on, describing the corresponding sample descriptives. We then set out the main results from our counterfactual simulations.

Simulation Procedure

The basic structure of the computation of a new equilibrium consists of a 'price' loop within a larger 'sociodemographics' loop. Having changed some primitive, we first calculate a new set of prices that clears the market. Here, Berry (1994) ensures that there is such a unique set of market-clearing prices (up to scale); in addition, Berry *et al*. (1995) provide a quick means of computing the market-clearing price vector.

We then take the new prices and the initial sociodemographic compositions of each neighborhood and go on to calculate the probability that each household chooses each housing type. Aggregating these choices to the neighborhood level, we also compute the corresponding predicted sociodemographic composition of each neighborhood. We then replace the initial neighborhood sociodemographic compositions with these new measures and start the price loop again, calculating a new set of market clearing prices given these updated neighborhood sociodemographic measures. We continue this process until the neighborhood sociodemographics converge.[12] The household location decisions corresponding to the final sociodemographic measures, along with the vector of housing prices that clears the market, then

---

strongly by households of this group (e.g., recent immigrants from China may have similar tastes for shops, restaurants, and other neighborhood amenities).

[12] In one variant of the simulation code, we allow school quality and crime rates to adjust in light of changing local demographics, allowing for further compounding. In the simulation results reported in this paper, we close down this channel of adjustment, potentially understating the effects of changes in primitives on stratification.

constitute the new equilibrium, based on which we can construct a variety of predicted segregation and consumption measures.

Sample and Pre-experiment Benchmark

As we have discussed above, equilibria are not computed during the estimation process. Not least, the class of models that ours falls into do not have generically unique equilibria. Thus it is unlikely that the prices and neighborhood compositions seen in the data constitute an equilibrium in the context of our sorting model, given the preference estimates we recover. For the purposes of carrying out simulations, it is therefore useful to start from a suitable benchmark in order to provide a comparison with subsequent counterfactuals.[13] To this end, we take the full set of model estimates and the simple structure of the equilibrium model, and slightly perturb tastes, say, over commuting. This sets in motion changes in market-clearing prices and neighborhood compositions until an equilibrium is found.

Based on this 'pre-experiment' equilibrium, it is straightforward to compute predicted measures of neighborhood stratification and the consumption of amenities for households with given sociodemographic characteristics; and these can be compared to the corresponding sample measures, as well as segregation and consumption measures predicted by the model directly, without imposing any equilibrium requirement.[14]

The presentation of the results from the 'pre-experiment' benchmark follows the same format as the results from the counterfactual simulations, so it is worth commenting on the general structure: for each equilibrium, we compute exposure rate measures of neighborhood segregation and predicted consumption measures for housing and neighborhood amenities, separately for households depending on their education and income.[15] We now describe the construction of these measures. For our education categorization, we assign households to three exhaustive, mutually exclusive categories, based on the educational attainment of the household head: high school or less, some college, and college degree or more. For income, we simply

[13] We note that multiplicity arises because of the presence of social interactions in the utility function – the potential dependence of household tastes on the characteristics of neighbors. None of the simulations we report involve shutting down these interactions, and so there is no guarantee that the counterfactual equilibria we describe are unique.
[14] We do not report segregation and consumption measures based on these simple model predictions here. They are very close indeed to the corresponding sample quantities, consistent with the close fit of the econometric model to the data.
[15] We also computed segregation as well as 'amenity' consumption measures on the basis of household education (whether the household head has a college degree or not) combined with an indicator denoting children present in the household – thus 2×2 categories.

assign households to one of the four income quartiles, where the relevant cutoffs are set based on the sample income distribution.

Exposure rates provide intuitive measures of the degree of neighborhood stratification (in this application, at the block group level) faced by households in a particular cell. Take, for instance, a given household in a given education cell. Based on the model predictions, the probability mass associated with that household is spread over housing units in the sample, and it is possible to construct the implied sociodemographic (in this case, educational) composition of each neighborhood that the household is associated with, having netted out the household's own presence. Averaging these predicted neighborhood compositions, using the household's probabilities as weights, it is possible to form the predicted average neighborhood compositions for the given household, based on the full set of choices in the choice set. In turn, we can average across all like households – for instance, households in the bottom education category – to determine the education exposure rates for households in that category. We construct analogous exposure rates for households with some college or a college degree, and separately, on the basis of household income.

For comparison, we also construct *sample* exposure rate measures, both by education and income. Here, for our sample education exposure rates, we simply record the block group neighborhood education composition for a household in a given category, then average the neighborhood compositions over all like households. Both for the pre-experiment and the sample, we obtain measures of the average neighborhood education compositions that households in each education category are exposed to, giving a total of nine exposure rates. These are arranged in columns of three, each column corresponding to the education compositions – either predicted or sample – for households falling into the education category denoted by the column heading. In the corresponding case of income, the relevant matrix of income quartile exposure rates – again, either predicted or sample – consists of 16 numbers, in four columns of four.

Table 5[16] shows education exposure rates for the sample (uppermost panel) and the pre-experiment predictions (bottom panel), with Table 6 showing corresponding income exposure rates. As guidance when reading the tables, the top panel of Table 5 shows sample education exposure rates. The top left-most entry, 0.431, indicates that in the sample, the typical household with no more than a high school education is found in a neighborhood (here, a Census block group) in which 43.1 percent of households has no more than a high school education. Looking down the first column, the remainder of the typical neighbors that less educated households are

[16] There is no Table 4 in this current draft!

exposed to consist of households with some college (22.5 percent), and households with a college degree (34.4 percent). The column entries necessarily sum to one.

The middle panel in the table gives the overall education distribution for the sample: 33.8 percent of households have at most a high school education, 22.3 percent have some college, and 43.8 percent have a college degree or more. It is clear that there is some degree of education segregation in the Bay Area sample, at least at the top and bottom ends of the education spectrum, with households tending to locate in neighborhoods where more households of their own education level are found, relative to the no-stratification case where the education composition of all neighborhoods mirrored the overall sample proportions. Thus, highly educated households are, on average, located in neighborhoods that are 51.8 percent highly educated (i.e. in the top education category), as opposed to 43.8 percent highly educated is there was 'even spreading.' This sample over-exposure comes at the expense of households in the two lower education categories, but especially households in the bottom education category (26.6 percent average exposure in contrast to the overall 33.8 percent). At the same time, it is worth noting that there is far from complete education stratification in the sample. Households in each education category tend to be found in neighborhoods that have reasonably high proportions of households in the other two categories.

A similar general pattern emerges for the sample income exposure rates. Looking at the top panel of Table 6, even income spreading would imply that all cell entries equaled 25 percent. Yet the largest numbers in each column are all found on the main diagonal of the 4-by-4 sample exposure rate matrix. And for households in the top income quartile, for instance, they are on average exposed to 37.2 rather than 25 percent highly educated neighbors; similar own-quartile over-exposure is apparent for households in the bottom quartile.

It is instructive to compare our sample exposure rates on the basis of education and income with the equilibrium exposure rates corresponding to a slight perturbation of household tastes. (As noted above, the straight predictions from the choice model, ignoring equilibrium considerations, are very close to the sample exposure rates.) In the case of education, given in the bottom panel of Table 5, the implied exposure rates exhibit somewhat greater 'own-type' stratification than the sample, especially at the top end. This may reflect the strong education interactions we estimate in the utility function, with highly educated households willing to pay relatively large amounts to live with similar highly educated households. Of note, this over-exposure pattern is *not* apparent for our income quartile exposure rate predictions. That is, the pre-experiment predictions for income are really rather similar to the sample, except at the very bottom end; and further, the predicted income segregation measures exhibit somewhat lower 'own-cell' stratification than the sample itself, rather than higher in the case of education.

We also report sample and predicted *consumption* measures based on the same household categorization: three education categories and four income quartile categories. The sample consumption measures, for housing (whether owned, and number of rooms) and local amenities (school quality and crime), are computed simply by averaging housing consumption and amenity levels over all households in a given cell, based on education or income quartile. In contrast, the predicted consumption measures use the equilibrium probabilities that spread households over the different choices in the choice set, from which we construct weighted-average consumption measures for households in a given education group or income quartile.

Considering first the sample consumption measures by education (see panel (1) of Table 9), ownership rates and houses sizes are all increasing in the education of the household head, as one might expect. For instance, while households with 'some college' have average ownership rates just over 55 percent, households with a college degree or more have ownership rates over 65 percent. There is a similar positive gradient in terms of school quality (as measured by average test scores), increasing from a score around 505 for the lowest to around 547 for the highest education group. And crime rates decline by about 35 percent on average, moving from the lowest to the most highly educated households.

The sample consumption measures by income quartile, shown in the top panel of Table 10, show a similar pattern, though the gradients tend to be steeper:[17] households in the top quartile have ownership rates of around 85 percent, contrasting sharply with households in the bottom quartile, whose ownership rates are well under half that (just over 37 percent); and house sizes are around 60 percent larger, comparing bottom to top quartile. Education exhibits a similar positive gradient, and the crime rate declines steeply, by almost 60 percent, moving from the bottom to the top income quartile.

Comparing sample with the pre-experiment consumption predictions – the top two panels of Table 9 (for education) and Table 10 (for income quartiles) – the fit tends to be close for the housing characteristics (ownership rates and number of rooms), reasonably close for test scores, and close for crime if we look at income quartiles. In terms of household education, the predicted crime pattern shows a slight upturn not apparent in the sample, and the predicted consumption profile for school quality is slightly flatter than in the sample.

Counterfactual Simulations – Commuting

To shed light on the relative importance of factors that may in practice contribute to residential mixing, we begin with two experiments that adjust tastes over commuting. In the econometric model, household preferences over commuting distance (the distance between each housing choice and the place of work of the household head) are allowed to vary with observable household characteristics, including their education, income, race, age, and family structure. The first experiment takes our taste estimates, which measure heterogeneous disutility of commuting, and reduces them by 50 percent. The second experiment is more extreme, switching off any disutility of commuting entirely. This second experiment is akin to removing geographic considerations from the decision process, aside from the pre-existing geographic variation in choice characteristics, including local public goods, across the urban area.

We start by examining counterfactual education exposure rates, comparing the predicted stratification from the two commuting experiments (panels (2) and (3) of Table 7) against the pre-experiment benchmark (top panel of Table 7). It is clear that education stratification increases in both of the counterfactual simulations relative to the pre-experiment case, evident in the increase in the entries on the main diagonals of the second and third panels. For instance, the exposure of the top education category increases by around 11 percent when commuting disutilities are cut in half, and a further 7 percent when they are switched off entirely. In terms of the magnitudes of the changes, proportionately larger effects tend to occur in moving from the pre-experiment to the counterfactual that cuts commuting disutilities by half, relative to the change between pre-experiment and switching off commuting disutilities entirely.

Considering counterfactual income exposure rates (see top three panels of Table 8), there is a similar qualitative pattern, though it is interesting to note that the changes in income stratification are more muted – for instance, own income segregation increasing by 5.4 and 7.6 percent, respectively for households in the top income quartile.

In terms of consumption patterns (see panels (2) – (4) of Tables 9 and 10), there are essentially no effects on consumption of housing characteristics by household education, though the consumption gradients for crime and school quality become somewhat steeper, with consumption of amenities rising slightly for the most highly educated households, falling slightly for the middle education group, and falling proportionately more for the bottom education category. This indicates that more highly educated households find it easier to sort into neighborhoods providing better amenities when commuting costs fall, though the reverse appears true for the lowest educated (perhaps because of relative fixity of supply). In terms of income,

[17] Part of this is mechanical, as there are four, not three categories.

again the commuting experiments have relatively minor effects on housing consumption, changes in consumption of school quality are barely perceptible, and there is only a very slight increase in the crime gradient, which again steepens. As one might expect, commuting distances do increase, by about 50 percent overall relative to the pre-experiment levels when commuting disutilities are halved, and by around 100 percent when they cease to matter at all.

In sum, the two commuting experiments indicate that reducing the importance of commuting in the household location decision results in monotonic but successively smaller increases in education and income stratification. Commuting distance increases across the board, and consumption of local amenities tends to increase for highly educated households and decline for the rest. Of note, even if commuting ceased to matter, our counterfactuals suggest that there would still be a good deal of within-neighborhood mixing on the basis of education and income, which is just to say that other factors in the model still serve to bring heterogeneous households together. Accordingly, our next experiment looks at the role of housing.

Counterfactual Simulations – Housing

Housing did not have an explicit role in Tiebout's original formulation, though subsequent research (see, for instance, the contributions of Oates and Hamilton) have made good any deficit in this regard. Analogous to the heterogeneous preferences over commuting distance, our sorting model interacts observable household characteristics with observed features of the housing stock (including the house price, whether owned, and house size). In our third counterfactual simulation, we switch off these heterogeneous housing preferences entirely. And again, we consider the associated effects on education and income stratification as well as household consumption patterns, disaggregated by household education category and income quartile.

In terms of stratification, this counterfactual leads to increases in education segregation that are remarkably similar to the effects of the first commuting experiment (that cut the disutility of commutes in half). This makes clear that the existing distribution of housing serves to integrate households on the basis of education to quite a sizeable degree. In marked contrast, switching off housing preferences leads to *reductions* in income stratification, producing reductions in own-quartile income segregation across the board. This shows that the differential affordability of larger housing serves to segregate households on the basis of income: our computations give us a sense of how much by. More generally, it indicates that income and education stratification do not work in an exactly co-linear way.

When we look at the effects of the housing experiment on consumption, switching off housing tastes leads, as one would anticipate, to a partial equalization of housing consumption (see panel (5) of Tables 9 and 10), especially on the basis of income. Similarly, there is a narrowing in the consumption of local amenities, with the gaps in the consumption of school quality and crime declining somewhat, especially across income quartiles.

Counterfactual Simulations – Commuting and Housing Combined

As a final experiment, we explore the consequences of switching off commuting preferences and preferences over housing at the same time. In terms of education stratification, the individual experiments are mutually reinforcing, as is apparent from panel (5) of Table 7; the combined experiment gives rise to the highest levels of education segregation relative to the pre-experiment benchmark. Even so, despite the radical nature of the experiment, education stratification is still far from complete: highly educated households still live in neighborhoods that also contain over 16 percent of households, on average, with no more than a high school education; and households in the bottom education category live in neighborhoods with 21 percent of households who are highly educated.

Consistent with the results from the housing experiment, the net effect on income segregation from the combined experiment looks remarkably like income segregation in the pre-experiment simulations, implying that the two experiments are almost entirely offsetting.

7 CONCLUSION

In this paper, we used a counterfactual equilibrium approach to shed light on factors that underpin the widespread sociodemographic integration observed in US cities. The approach allows us to change primitives of an econometric sorting model, estimated using very rich Census data, then compute the implied stratification associated with the counterfactual equilibria that arise.

Our first simulations explored the role of changing the disutility of commutes counterfactually. As commuting matters less to households, so residential stratification at the Census block level increases markedly, especially in terms of education. This makes clear that employment geography and non-trivial commuting costs serve to bring heterogeneous households together in the same residential neighborhoods. As has been noted by Oates and Schwab (among others), firm production typically requires complementary labor inputs, combining low and high skill workers, and this serves as an integrating force. Never-the-less, even when commuting considerations are switched off entirely, a fairly high degree of residential mixing on the basis of education and income persists. When we look at implied consumption of amenities, our results suggest, as is plausible, that as distance to work matters less, so stratification on the basis of local public goods consumption increases: highly educated households tend to locate in neighborhoods with better schools and lower crime than in the benchmark pre-experiment equilibrium, while the reverse is true for less-educated households.

Examining the separate contribution of heterogeneous tastes for housing, our findings point to an interesting contrast: while switching off tastes for housing leads to quite marked increases in education stratification, income stratification declines, pointing to the fact that the distribution of housing (with larger houses clustering in more affluent neighborhoods) serves to segregate households on the basis of income. When housing characteristics cease to matter directly in the location decision, we see some equalization in the consumption of amenities (in addition to the anticipated equalization in housing consumption).

Our approach takes as given the distribution of employment and the characteristics of housing. A long line of research by Anas and coauthors has sought to endogenize firm locations, which is a natural if technically challenging avenue to take the literature in. And several recent papers have developed sophisticated econometric models of housing supply, building on earlier work by Henderson and others emphasizing the role of developers. Dynamic equilibrium models that incorporate housing, amenable to counterfactual simulation analysis, seem highly appealing. As yet, they remain beyond the limits of the existing literature.

REFERENCES

Anas, Alex, (1982), *Residential Location Markets and Urban Transportation: Economic Theory, Econometrics and Public Policy Analysis*, Academic Press, New York.

_______ (1990), "Taste Heterogeneity and Urban Spatial Structure: Monocentric Analysis and the Logit Model Reconciled," *Journal of Urban Economics*, 18, 7-35.

_______ (2002), "Prejudice, Exclusion, and Compensating Transfers: The Economics of Ethnic Segregation," *Journal of Urban Economics*, 52, 409-32.

Anas, Alex and Ikki Kim, (1996), "General Equilibrium Models of Polycentric Urban Land Use with Endogenous Congestion and Job Agglomeration," *Journal of Urban Economics*, 40, 232-56.

Anas, Alex and Richard Arnott, (1993), "Development and Testing of the Chicago Prototype Housing Market Model," *Journal of Housing Research,* 4(1): 73-129.

_______, (1997), "Taxes and Allowances in a Dynamic Equilibrium Model of Urban Housing with a Size-Quality Hierarchy," *Regional Science and Urban Economics*, 27, 547-80.

Bajari, Patrick, and Matthew Kahn (2005), "Estimating Housing Demand with an Application to Explaining Racial Segregation in Cities," *Journal of Business and Economic Statistics*, 23(1): 20-33.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan (2007), "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 115(4): 588-638.

Bayer, Patrick, Robert McMillan, and Kim Rueben (2009), "An Equilibrium Model of an Urban Housing Market," revised version of NBER Working Paper 10865.

Benabou, Roland (1993), "The Workings of a City: Location, Education, and Production," *Quarterly Journal of Economics*, 108(3), pp.619-652.

________ (1996), "Heterogeneity, Stratification, and Growth: Macroeconomic Implications of Community Structure and School Finance," *American Economic Review*, 86(3), 584-609.

Berry, Steven (1994), "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, Vol. 25, pp. 242-262.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica*, 63: 841-890.

Bewley, Truman (1981), "A Critique of Tiebout's Theory of Local Public Expenditures," *Econometrica*, 49: 713-740.

Black, Sandra (1999), "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics*, May.

Cullen, Julie, Brian Jacob, and Steven Levitt (2006), "The Effect of School Choice on Student Outcomes: Evidence from Randomized Lotteries," *Econometrica*, 74(5): 1191-1230.

Cutler, David, and Edward Glaeser (1997), "Are Ghettos Good or Bad?" *Quarterly Journal of Economics*, 112(3): 827-882.

Debreu, Gerard (1969), "Economies with a Finite Set of Equilibria," Irving Fisher lecture at the Brussels meeting of the Econometric Society, September, reprinted in *Mathematical Economics: Twenty Papers of Gerard Debreu*, Cambridge University Press, 1983.

Epple, Dennis, Radu Filimon, and Thomas Romer (1984), "Equilibrium Among Local Jurisdictions: Towards an Integrated Approach of Voting and Residential Choice," *Journal of Public Economics*, 24: 281-304.

________ (1993), "Existence of Voting and Housing Equilibrium in a System of Communities with Property Taxes," *Regional Science and Urban Economics*, 23: 585-610.

Epple, Dennis, and Glenn Platt (1998), "Equilibrium and Local Redistribution in an Urban Economy when Households Differ in Both Preferences and income," *Journal of Urban Economics*, 43: 23-51.

Epple, Dennis and Richard Romano (1996), "The Public Provision of Private Goods," *Journal of Political Economy*, 104(1): 57-84.

Epple, Dennis and Richard Romano (1998), "Competition Between Private and Public Schools, Vouchers, and Peer Group Effects," *American Economic Review*, 88(1): 33-62.

Epple, Dennis, Thomas Romer and Holger Sieg (2001), "Interjurisdictional Sorting and Majority Rule: An Empirical Analysis," *Econometrica*, 69(6): 1437-1455.

Epple, Dennis and Holger Sieg (1999), "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy*, 107(4): 645-681.

Epple, Dennis, and Allan Zelenitz (1981), "The Implications of Competition among Jurisdictions: Does Tiebout Need Politics?" *Journal of Political Economy*, 89(4): 1197-1218.

Fernandez, Raquel and Richard Rogerson (1996), "Income Distribution, Communities, and the Quality of Public Education." *Quarterly Journal of Economics*, 111(1): 135-164.

________ (2003), "Equity and Resources: An Analysis of Educational Finance Systems," *Journal of Political Economy*, 111(4).

Ferreyra, Maria (2007), "Estimating the Effects of Private School Vouchers in Multi-District Economies," *American Economic Review*, 97(3): 789-817.

McFadden, Daniel (1978), "Modeling the Choice of Residential Location," in eds. Karlquist, A. *et al.*, Spatial Interaction Theory and Planning Models, Elsevier North-Holland, New York.

Nechyba, Thomas J. (1999), "School Finance-Induced Migration and Stratification Patterns: the Impact of Private School Vouchers," *Journal of Public Economic Theory*, Vol. 1.

________ (2000), "Mobility, Targeting and Private School Vouchers," *American Economic Review* 90(1), 130-46.

Nechyba, Thomas J., and Robert P. Strauss, (1998), "Community Choice and Local Public Services: A Discrete Choice Approach," *Regional Science and Urban Economics*, Vol. 28, 51-73.

Quigley, John M., (1985), "Consumer Choice of Dwelling, Neighborhood, and Public Services," *Regional Science and Urban Economics*, Vol. 15(1).

Sieg, Holger, V. Kerry Smith, H. Spencer Banzaf and Randall Walsh (2004), "Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods," *International Economic Review*, 45(4): 1047-1077.

Tiebout, Charles M. (1956), "A Pure Theory of Local Expenditures," *Journal of Political Economy*, 64: 416-424.

Table 1. Sample Statistics Comparing the Full Sample with Houses within 0.20 miles of a Boundary

Sample	**full sample**		**within 0.20 miles of boundaries**				
			boundary sample	**high test score side**	**low test score side**	**difference in means**	**test of difference**
Observations	**242,100**		**27,548**	**13,612**	**13,936**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mean	S.D.	Mean	Mean	Mean	((4) - (5))	t-statistic
Housing Prices							
House value (if owned)	297,700	178,479	250,005	259,475	240,756	18,719	4.15
Monthly rent (if rented)	744	316	678	688	669	18.80	1.73
School Quality							
Average test score	527	74	507	544	471	74	25.44
Housing Characteristics							
1 if unit owned	0.60	0.49	0.54	0.55	0.53	0.02	0.89
Number of rooms	5.11	1.99	4.96	5.02	4.90	0.12	1.56
1 if built in 1980s	0.14	0.35	0.11	0.11	0.11	0.00	-0.31
1 if built in 1960s or 1970s	0.39	0.49	0.34	0.35	0.33	0.01	0.84
Elevation	210	179	176	178	173	6	1.64
Population density	0.43	0.50	0.39	0.38	0.40	-0.02	-1.38
Neighborhood Sociodemographics							
% Census block group white	0.68	0.23	0.61	0.63	0.60	0.03	3.40
% Census block group black	0.08	0.16	0.18	0.17	0.20	-0.03	-3.15
% Census block group coll deg or more	0.44	0.20	0.41	0.44	0.39	0.05	6.18
Average block group income	54,742	26,075	46,271	47,718	44,857	2,861	2.61

Note: This table reports summary statistics for the key variables included in the analysis. The boundary sample includes all houses located within 0.20 miles of a boundary with another school attendance zone. A house is considered to be on the 'high' ('low') side of a boundary if the test score at its local school is greater (less) than the corresponding test score for the closest house on the opposite side of an attendance zone boundary. Sample statistics are reported for the high- and low-side of boundaries for which the test score gap is in excess of the median gap (38.4 points) in columns (4) and (5), respectively. Column (7) reports the t-statistic for a test of the hypothesis that the mean of the variable listed in the row heading does not vary across school attendance zone boundaries. This test conditions on boundary fixed effects (so as to compare houses on opposite sides of the same boundary) and adjusts for the clustering of observations at the Census block group level.

Table 2: Delta Regressions - Implied Mean Willingness to Pay

Sample	**Within 0.20 Miles of Boundary**	
Observations	**27,458**	
Boundary Fixed Effects	**No**	**Yes**
Number of Rooms	91.7	91.5
	(7.1)	(13.9)
Built in 1980s	92.0	95.4
	(9.8)	(15.1)
Built in 1960s/70s	9.2	7.3
	(3.3)	(2.4)
Owner-occupied	69.1	51.0
	(4.7)	(6.1)
Average Test Score (in standard deviations)	18.0	19.7
	(8.3)	(7.4)
% Neighborhood Black	-404.8	-104.8
	(41.4)	(36.9)
% Neighborhood Hispanic	-88.4	-3.5
	(32.5)	(31.0)
% Neighborhood Asian	-39.7	-5.3
	(30.2)	(32.4)
% Neighborhood College Degree or more	183.5	104.6
	(26.4)	(31.8)
Average Neighborhood Income (/10000)	30.7	36.3
	(3.7)	(6.6)
Distance to Work	-25.6	-52.2
	(0.3)	(0.5)

Note: All regressions shown in the table also include controls for whether house was built in 1960-1979, elevation, population density, crime, land use (% industrial, % residential, % commercial, % open space, % other) in 1, 2 and 3 mile rings around each location. The dependent variable is the monthly user cost of housing, which equals monthly rent for renter-occupied units and a monthly user cost for owner-occupied housing, calculated as described in the text. Standard errors corrected for clustering at the school level are reported in parentheses.

Table 3. Heterogeneity in Marginal Willingness to Pay for Housing and Neighborhood Characteristics

	House Characteristics			**Neighborhood Characteristics**					
	Own vs. Rent	+1 Room	Built in 1980s vs. pre-1960	+10% Black vs. White	+10% Hisp vs. White	+10% Asian vs. White	+10% College Educated	Blk Group Avg Income + $10,000	Average Test Score +1 s.d.
Mean MWTP	51.0	91.5	95.4	-10.5	-0.4	-0.5	10.5	36.3	19.7
	(6.1)	(13.9)	(15.1)	(3.7)	(3.1)	(3.2)	(3.2)	(6.6)	(7.4)
Household Income (+$10,000)	21.7	4.8	9.3	-1.2	0.8	0.1	1.4	0.9	1.4
	(0.7)	(0.2)	(0.8)	(0.4)	(0.4)	(0.3)	(0.2)	(0.1)	(0.3)
Children Under 18 vs.	-13.3	37.5	-26.0	11.9	17.2	12.6	-16.1	2.4	7.4
No Children	(7.0)	(1.8)	(8.5)	(3.0)	(3.3)	(2.7)	(2.2)	(1.2)	(3.6)
Black vs. White	-67.7	3.8	5.4	98.3	46.7	48.3	18.4	-1.2	-14.3
	(13.0)	(3.5)	(17.4)	(3.9)	(5.6)	(5.1)	(4.5)	(2.2)	(7.4)
Hispanic vs. White	-8.2	-14.5	-6.3	30.9	85.6	18.0	6.3	1.1	-4.1
	(10.1)	(2.7)	(12.2)	(3.9)	(4.0)	(4.2)	(3.4)	(1.4)	(6.0)
Asian vs. White	113.5	-22.1	43.9	28.1	22.3	95.2	0.4	0.7	7.0
	(9.6)	(2.2)	(11.3)	(3.8)	(4.4)	(3.5)	(2.6)	(1.5)	(5.5)
College Degree or More vs.	37.6	-0.6	40.2	9.2	-4.6	-13.4	58.0	0.3	13.0
Some College or Less	(8.1)	(2.2)	(9.7)	(3.1)	(3.9)	(3.0)	(2.3)	(1.4)	(3.6)

Note: The first row of the table reports the mean marginal willingness-to-pay for the change reported in the column heading. The remaining rows report the difference in willingness to pay associated with the change listed in the row heading, holding all other factors equal. The full heterogeneous choice model includes 135 interactions between nine household characteristics and fifteen housing and neighborhood characteristics. The included household characteristics are household income, the presence of children under 18, and the race/ethnicity (Asian, black, Hispanic, white), educational attainment (some college, college degree or more), work status, and age of the household head. The housing and neighborhood characteristics are the monthly user cost of housing, distance to work, average test score, whether the house is owner-occupied, number of rooms, year built (1980s, 1960-1979, pre-1960), elevation, population density, crime, and the racial composition (% Asian, % black, % Hispanic, % white) and average education (% college degree) and household income for the corresponding Census block group. Standard errors are reported in parentheses.

Table 5: Education Exposure Rates - Sample and Pre-Experiment

Sample Exposure Rates

	Household Education Level		
	≤ High school	Some college	≥ College degree
Neighborhood percent ≤ High school	0.431	0.341	0.266
Neighborhood percent with Some college	0.225	0.235	0.216
Neighborhood percent ≥ College degree	0.344	0.424	0.518
Total	1.000	1.000	1.000

Sample Education Distribution

Household Education Level		
≤ High school	Some college	≥ College degree
0.338	0.223	0.438

Pre-experiment Simulation Exposure Rates

	Household Education Level		
	≤ High school	Some college	≥ College degree
Neighborhood percent ≤ High school	0.449	0.367	0.238
Neighborhood percent with Some college	0.243	0.245	0.198
Neighborhood percent ≥ College degree	0.308	0.388	0.564
Total	1.000	1.000	1.000

Notes:

Each column gives the average exposure of households whose type is given by the relevant column heading to neighbors in the row category.
Neighborhoods are defined at the block group level.

Table 6: Income Exposure Rates - Sample and Pre-experiment

Sample Exposure Rates

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Neighborhood percent Income Quartile 1	0.351	0.271	0.218	0.164
Neighborhood percent Income Quartile 2	0.268	0.271	0.250	0.206
Neighborhood percent Income Quartile 3	0.217	0.252	0.273	0.258
Neighborhood percent Income Quartile 4	0.163	0.207	0.258	0.372
	1.000	1.000	1.000	1.000

Pre-experiment Simulation Exposure Rates

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Neighborhood percent Income Quartile 1	0.319	0.278	0.239	0.169
Neighborhood percent Income Quartile 2	0.275	0.265	0.249	0.206
Neighborhood percent Income Quartile 3	0.238	0.250	0.257	0.256
Neighborhood percent Income Quartile 4	0.168	0.207	0.255	0.369
	1.000	1.000	1.000	1.000

Notes

Each column gives the average exposure of households whose type is given by the relevant column heading to neighbors in the row category.

Neighborhoods are defined at the block group level.

Table 7: Education Exposure Rates - Simulation Results

(1) Pre-experiment

	Household Education Level		
	≤ High school	Some college	≥ College degree
Neighborhood percent ≤ High school	0.449	0.367	0.238
Neighborhood percent with Some college	0.243	0.245	0.198
Neighborhood percent ≥ College degree	0.308	0.388	0.564

(2) Counterfactual: ***cut disutility of commutes in half***

	Household Education Level		
	≤ High school	Some college	≥ College degree
Neighborhood percent ≤ High school	0.485	0.392	0.197
Neighborhood percent with Some college	0.259	0.259	0.178
Neighborhood percent ≥ College degree	0.256	0.349	0.624

(3) Counterfactual: ***switch off disutility of commutes***

	Household Education Level		
	≤ High school	Some college	≥ College degree
Neighborhood percent ≤ High school	0.503	0.410	0.174
Neighborhood percent with Some college	0.271	0.273	0.161
Neighborhood percent ≥ College degree	0.226	0.316	0.664

(4) Counterfactual: ***switch off housing preferences***

	Household Education Level		
	≤ High school	Some college	≥ College degree
Neighborhood percent ≤ High school	0.478	0.396	0.202
Neighborhood percent with Some college	0.261	0.261	0.175
Neighborhood percent ≥ College degree	0.261	0.344	0.623

(5) Counterfactual: ***switch off disutility of commutes and housing preferences***

	Household Education Level		
	≤ High school	Some college	≥ College degree
Neighborhood percent ≤ High school	0.507	0.428	0.162
Neighborhood percent with Some college	0.282	0.288	0.145
Neighborhood percent ≥ College degree	0.210	0.285	0.693

Notes:
Each column gives the average exposure of households whose type is given by the relevant column heading to neighbors in the row category.
All the columns sum to 1.
Neighborhoods are defined at the block group level.

Table 8: Income Exposure Rates - Simulation results

(1) Pre-experiment

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Neighborhood percent Income Quartile 1	0.319	0.278	0.239	0.169
Neighborhood percent Income Quartile 2	0.275	0.265	0.249	0.206
Neighborhood percent Income Quartile 3	0.238	0.250	0.257	0.256
Neighborhood percent Income Quartile 4	0.168	0.207	0.255	0.369

(2) Counterfactual: ***cut disutility of commutes in half***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Neighborhood percent Income Quartile 1	0.326	0.281	0.238	0.159
Neighborhood percent Income Quartile 2	0.278	0.268	0.250	0.199
Neighborhood percent Income Quartile 3	0.237	0.251	0.260	0.253
Neighborhood percent Income Quartile 4	0.158	0.199	0.253	0.389

(3) Counterfactual: ***switch off disutility of commutes***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Neighborhood percent Income Quartile 1	0.332	0.282	0.237	0.154
Neighborhood percent Income Quartile 2	0.279	0.269	0.250	0.196
Neighborhood percent Income Quartile 3	0.236	0.252	0.261	0.253
Neighborhood percent Income Quartile 4	0.153	0.197	0.252	0.397

(4) Counterfactual: ***switch off housing preferences***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Neighborhood percent Income Quartile 1	0.306	0.270	0.242	0.187
Neighborhood percent Income Quartile 2	0.267	0.259	0.249	0.220
Neighborhood percent Income Quartile 3	0.241	0.250	0.256	0.254
Neighborhood percent Income Quartile 4	0.186	0.220	0.254	0.339

(5) Counterfactual: ***switch off disutility of commutes and housing preferences***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Neighborhood percent Income Quartile 1	0.314	0.274	0.242	0.175
Neighborhood percent Income Quartile 2	0.271	0.263	0.251	0.211
Neighborhood percent Income Quartile 3	0.241	0.252	0.257	0.251
Neighborhood percent Income Quartile 4	0.174	0.211	0.250	0.364

Notes

Each column gives the average exposure of households whose type is given by the relevant column heading to neighbors in the row category.

All the columns sum to 1.

Neighborhoods are defined at the block group level.

Table 9: Consumption Rates by Education Category

(1) Sample

	Household Education Level		
	≤ High school	Some college	≥ College degree
ownership rate	0.549	0.555	0.655
number of rooms	4.7	5.0	5.5
school test score	504.7	525.5	546.8
crime rate	10.4	7.5	6.8

Simulation results:
(2) Pre-experiment

	Household Education Level		
	≤ High school	Some college	≥ College degree
ownership rate	0.547	0.561	0.653
number of rooms	4.73	5.10	5.42
school test score	508.54	529.60	541.77
crime rate	9.75	7.16	7.50

(3) Counterfactual: ***cut disutility of commutes in half***

	Household Education Level		
	≤ High school	Some college	≥ College degree
ownership rate	0.552	0.562	0.649
number of rooms	4.70	5.08	5.45
school test score	504.79	527.94	545.42
crime rate	10.96	7.44	6.42

(4) Counterfactual: ***switch off disutility of commutes***

	Household Education Level		
	≤ High school	Some college	≥ College degree
ownership rate	0.550	0.564	0.650
number of rooms	4.69	5.09	5.45
school test score	504.41	527.74	545.92
crime rate	11.49	7.65	5.89

(5) Counterfactual: ***switch off housing preferences***

	Household Education Level		
	≤ High school	Some college	≥ College degree
ownership rate	0.579	0.616	0.601
number of rooms	4.91	5.21	5.23
school test score	510.89	529.91	539.84
crime rate	9.38	6.97	7.86

(6) Counterfactual: ***switch off disutility of commutes and housing preferences***

	Household Education Level		
	≤ High school	Some college	≥ College degree
ownership rate	0.534	0.604	0.642
number of rooms	4.70	5.14	5.42
school test score	505.73	527.56	544.96
crime rate	11.10	7.43	6.33

Note: Table gives average consumption levels of the row characteristic for households in the column heading category.

Table 10: Consumption Rates by Income Quartile

(1) Sample

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ownership rate	0.373	0.499	0.671	0.847
number of rooms	3.9	4.6	5.4	6.5
school test score	504.2	518.5	530.0	558.6
crime rate	12.4	8.6	6.7	5.0

Simulation results
(2) Pre-experiment

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ownership rate	0.420	0.491	0.625	0.852
number of rooms	4.21	4.62	5.24	6.38
school test score	506.43	518.09	530.09	556.69
crime rate	11.52	9.09	7.19	4.93

(3) Counterfactual: ***cut disutility of commutes in half***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ownership rate	0.416	0.485	0.625	0.862
number of rooms	4.18	4.59	5.24	6.45
school test score	505.32	517.29	529.91	558.63
crime rate	12.09	9.14	7.05	4.45

(4) Counterfactual: ***switch off disutility of commutes***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ownership rate	0.410	0.483	0.627	0.867
number of rooms	4.16	4.59	5.24	6.48
school test score	505.08	517.57	530.23	558.45
crime rate	12.39	9.12	6.92	4.29

(5) Counterfactual: ***switch off housing preferences***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ownership rate	0.539	0.568	0.603	0.679
number of rooms	4.72	4.91	5.15	5.68
school test score	508.86	519.99	530.15	552.37
crime rate	10.56	8.66	7.44	6.04

(6) Counterfactual: ***switch off disutility of commutes and housing preferences***

	Household Income			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ownership rate	0.510	0.559	0.604	0.716
number of rooms	4.57	4.85	5.15	5.88
school test score	506.59	518.52	529.62	556.53
crime rate	11.59	8.77	7.23	5.16

Note: Table gives average consumption levels of the row characteristic for households in the column heading category.