

TIEBOUT SORTING, SOCIAL MULTIPLIERS AND THE DEMAND FOR SCHOOL QUALITY*

Patrick Bayer
Department of Economics
Yale University

Fernando Vendramel Ferreira
The Wharton School
University of Pennsylvania

Robert McMillan
Department of Economics
University of Toronto

October 2004

Abstract

In many theoretical public finance models, school quality plays a central role as a determinant of household location choices and in turn, of neighborhood stratification. In contrast, the recent empirical literature has almost universally concluded that the direct effect of school quality on housing demand is weak, a conclusion that is robust across a variety of research designs. Using an equilibrium model of residential sorting, this paper closes the gap between these literatures, providing clear evidence that the full effect of school quality on residential sorting is significantly larger than the direct effect – four times as great for education stratification, twice for income stratification. This is due to a strong social multiplier associated with heterogeneous preferences for peers and neighbors; initial changes in school quality set in motion a process of re-sorting on the basis of neighborhood characteristics that reinforces itself, giving rise to substantially larger stratification effects.

Key words: Tiebout sorting, neighborhood stratification, social multiplier, school quality, education demand, capitalization, neighborhood choice, general equilibrium.

JEL: I20, H41, R21

* We are grateful to Pat Bajari, Steve Berry, Sandra Black, David Card, Ken Chay, Dennis Epple, Caroline Hoxby, Brian Jacob, David Lee, Tom Nechyba, Jesse Rothstein, Kim Rueben, Holger Sieg, Doug Staiger, and Chris Timmins for valuable discussions. We also thank Pedro Cerdan and Jackie Chou for help in assembling the data. This research was conducted at the California Census Research Data Center: our thanks to the CCRDC, and to Ritch Milby in particular. Financial support for this project provided by the NSF under grant SES-0137289, the Public Policy Institute of California, SSHRC Canada, and CAPES – Brazil is gratefully acknowledged. Correspondence may be sent to: Patrick Bayer, 37 Hillhouse Avenue, Dept. of Economics, Yale University, New Haven, CT 06511, patrick.bayer@yale.edu, (203) 432-6292.

1 INTRODUCTION

The notion that local public goods influence household location decisions has been a key idea in local public finance at least since Tiebout's seminal 1956 paper. According to Tiebout, variation in packages of local public goods leads households to "vote with their feet," providing a primary driving force behind household sorting across jurisdictions and neighborhood socioeconomic stratification. This idea has been central to a more recent theoretical literature examining the effects of education policy in the presence of household sorting.¹ In that literature, changes in school quality often have striking effects on neighborhood stratification as households resort across neighborhoods in response. Further, equilibrium models of residential sorting prove particularly useful in tracing many complex and otherwise difficult-to-predict effects of policy:² effects appear in general equilibrium that are often not apparent from a partial equilibrium perspective.³

At the same time, a growing recent empirical literature has almost universally concluded that the effect of school quality on the demand for schools and their corresponding neighborhoods is quite small. Strikingly, this conclusion is robust across a wide variety of research designs including: (i) examining differences in housing prices along school assignment boundaries or as school assignments are changed over time;⁴ (ii)

¹ Important examples include work by Benabou (1993, 1996), Fernandez and Rogerson (1996, 1998), and Nechyba (1999, 2000). Foundational papers by Epple, Filimon and Romer (1984, 1993) focus on conditions needed to prove existence of equilibrium in multi-community models that incorporate voting.

² Influential research by Epple and Romano (1998) uses an equilibrium approach to study education market equilibria that involve households choosing between public and private schools. While the housing market is abstracted from, this work has proved useful in an empirical context, generating predictions that are largely supported by the data - see Epple, Figlio and Romano (2004).

³ The work of Fernandez and Rogerson (1996) and Nechyba (2000) graphically demonstrates the importance of accounting for household mobility when assessing the effects of education policy.

⁴ Black (1999) pioneered the use of school assignment boundary fixed effects, estimating that a school-level standard deviation in average test score is associated with approximately a 2 percent increase in house

estimating the heterogeneity in willingness to pay for school quality using a model of residential sorting;⁵ (iii) examining the academic performance of winners versus losers following randomized school choice lotteries;⁶ and (iv) examining how student performance varies with measures of school district enrollment concentration across metropolitan areas.⁷ Thus taken at face value, there appears to be a gap between the central conclusion of the recent empirical literature and the emphasis placed on school-related sorting in the long theoretical tradition in local public finance.

This gap may not be as large as it appears at first glance. The focus of most of the empirical papers has been to isolate the *direct* effect of school quality on demand by controlling carefully for the correlation of school quality with the characteristics of peers and neighbors as well as other unobserved neighborhood attributes. In so doing, these same studies have also made it clear that households place a high value directly on the characteristics of peers and neighbors.⁸ The existence of such strong preferences for peers and neighbors, however, raises the possibility that a change in school quality may also have a sizeable *indirect* effect on demand as households move in response to any changes in peer and neighbor characteristics induced by the initial change in school quality. The full effect of a change in school quality on demand will be larger than the

value. Kane, Staiger, and Riegg (2004) report almost an identical result using boundary fixed effects. This latter paper also reports results based on changes in school assignment, although these are less conclusive.

⁵ Using the methodology presented in this paper, Bayer, Ferreira, and McMillan (2003) incorporate school district boundary fixed effects in estimating a model of residential sorting, returning an estimated mean marginal willingness to pay for a standard deviation increase in average test score of approximately 2 percent of house value; they also find evidence of heterogeneity around this mean, as well as strong preferences relating to the characteristics of neighbors.

⁶ Cullen, Jacob, and Levitt (2003) find little evidence that winners of randomized lotteries perform better in the schools that they subsequently select than losers who do not have the same degree of choice.

⁷ Rothstein (2003) finds little evidence that sorting is related to school rather than peer characteristics.

⁸ For example, in most of the empirical studies just referenced, estimates of the demand for school quality fall substantially when the authors control carefully for the sociodemographic characteristics of peers or neighbors. This suggests that much of what initially appears to be evidence of strong demand for school quality is more properly attributed to the demand for the characteristics of peers or neighbors.

direct effect identified in the literature when (i) there is considerable heterogeneity in the demand for school quality, and (ii) households have strong preferences for the characteristics of their neighbors, leading the initial sorting to be reinforced.

Given the possibility of a strong social multiplier of this type, the goal of this paper is to re-examine estimates of the demand for school quality through the lens of an equilibrium model of residential sorting. In particular, we present estimates of the effect of school-related sorting on neighborhood stratification by income and education, comparing a partial equilibrium measure that accounts for only the direct effect of school quality on stratification with a general equilibrium measure that also permits additional resorting as neighbors change. This exercise also informs our understanding of the nature of demand in the education marketplace more generally, providing for example a clear indication of the strength of market forces. If the estimated demand for school quality were as inelastic as the direct effects reported in the empirical literature suggest, one would certainly conclude that Tiebout-type sorting provides at most limited competitive pressure for schools, as a sizeable improvement in quality on the part of a school would have only a small effect on its demand. If the general equilibrium impact of a change in school quality on sorting is substantially larger than this direct effect, however, the competitive pressure induced by Tiebout-type sorting would be greater as well.⁹

To estimate the importance of school-related sorting on residential stratification, we follow the methodology developed in Bayer, Ferreira and McMillan (2003).¹⁰ In particular, we estimate a model of residential sorting using detailed Census data on a

⁹ Hoxby (2000) indicates that Tiebout sorting has a substantial impact on the competitiveness of the education marketplace, with metropolitan areas that afford more Tiebout choice having more productive public schools.

large (1-in-7) sample of households in the San Francisco Bay Area in 1990.¹¹ These data, available on a restricted basis at Census Research Data Centers throughout the country, specify the Census block in which each household lives, thereby characterizing each household's residential location very precisely both in terms of the characteristics of its neighbors and, through data from outside sources, other neighborhood attributes including information related to schools and school assignment.

The availability of individual-level data linked to such precise geographic indicators makes it possible to estimate a very flexible model of residential sorting,¹² both in terms of the structure of the model itself and in the way we use the data to estimate it. Our empirical specification allows household location decisions to be influenced by a large number of factors – school quality, consumption, housing quality and quantity, fixed neighborhood attributes, endogenous characteristics of neighbors, distance to work – and allows the demand for each of these factors to be affected by a household's own characteristics, including education, income, race, and family structure. The data also make it possible to incorporate the empirical strategy of examining variation along school district boundaries extending the work of Black (1999); as we show below, this permits the estimation of the full set of the model's parameters in a manner that deals with an important and pervasive endogeneity problem associated with the correlation of both

¹⁰ This previous paper sets out the model and estimation procedure in detail, describing, for example, the relationship between the identification of the broader sorting model used here and the hedonic price regressions commonly estimated in this literature.

¹¹ Given California's finance system, in which local jurisdictions have almost no control over the level of local taxation or spending on public goods including schools, this should provide a lower bound on the importance of school-related sorting in other states.

¹² A number of recent papers have estimated equilibrium sorting models. The first paper in this new empirical literature, by Epple and Sieg (1999), develops an estimator for the sorting model of Epple, Filimon and Romer (1984, 1993), used for example by Sieg *et al.* (2004) to explore the general equilibrium impacts of air quality improvements in the Los Angeles Basin. See Ferreyra (2003) and Walsh (2004), Bayer, McMillan, and Rueben (2004b) and Timmins (2003) for other applications using equilibrium sorting models.

school quality and neighborhood sociodemographic characteristics with unobserved neighborhood characteristics.

In line with other estimates in the literature, our estimate of the mean marginal willingness to pay for school quality is reasonably small, on the order of 2 percent of house value for a one standard deviation increase in school quality. An important feature of our approach is that it also permits the recovery of a broad set of preference parameters that are heterogeneous across different types of household, revealing that households have strong and very heterogeneous preferences for the characteristics of their neighbors. When combined with the underlying heterogeneity in the demand for school quality, these strong social interactions in the location decision are exactly the sort of preferences under which the general equilibrium effect of an exogenous change in school quality on neighborhood stratification can be substantially greater than the partial equilibrium effect.

Given these parameter estimates, the primary economic analysis of the paper quantifies the impact of school-related sorting on neighborhood stratification. Here, we use an intuitive simulation approach, solving for a new equilibrium in a counterfactual environment that eliminates household preferences for school quality from the location decision. This provides a characterization of the way the metropolitan area would look in a world in which sorting on the basis of public school quality played no role whatsoever. Our general equilibrium results show that school quality is an important driving force behind stratification: if school quality were eliminated as a determinant of household locations, stratification would decline by over 30 percent on the basis of education and by around 25 percent on the basis of income. These findings contrast with the partial equilibrium predictions of the model that consider only the direct effect of eliminating

school quality from the location decision, not allowing the subsequent changes in neighborhood sociodemographic compositions to feed back into household utility functions. Such partial equilibrium predictions give rise to only modest reductions in education and income stratification – of the order of 8-10 percent. Taken together, these findings are supportive of the earlier applied theory literature: changes in school quality can be expected to have significant effects on neighborhood stratification in equilibrium, effects that are much larger than the estimated direct effects of demand for school quality might suggest.

The rest of the paper is organized as follows: a model of residential sorting is presented in the next section. Section 3 discusses the estimation procedure. Section 4 describes the restricted-access Census micro data, Section 5 presents estimation results, and Section 6 presents the general equilibrium simulation results. Section 7 concludes.

2 A MODEL OF RESIDENTIAL SORTING

This section describes an equilibrium model of a self-contained urban housing market in which households sort themselves among the set of available housing types and locations.¹³ The model consists of two key elements: the household residential location decision problem and a market-clearing condition. While it has a simple structure, the model allows households to have heterogeneous preferences defined over housing and neighborhood attributes in a very flexible way; it also allows for housing prices and neighborhood sociodemographic compositions to be determined in equilibrium.

¹³ More details concerning the model described here can be found in Bayer, Ferreira, and McMillan (2003) and Bayer, McMillan, and Rueben (2004a).

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. The utility function specification is based on the random utility model developed in McFadden (1973, 1978) and the specification of Berry, Levinsohn, and Pakes (1995), which includes choice-specific unobservable characteristics. 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, 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 :

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

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, ζ_h , and an individual-specific term, ε_h^i . A useful interpretation of ζ_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_j^i , for $j \in \{X, Z, d, p\}$, varies with a household's own characteristics according to:

$$(2) \quad \alpha_j^i = \alpha_{0j} + \sum_{r=1}^R \alpha_{rj} z_r^i,$$

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 households have preferences defined distinctly over each choice characteristic including both school quality and neighborhood sociodemographic characteristics.¹⁴ 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.¹⁵ The additional flexibility of the horizontal model is especially relevant for this paper as it is the magnitude of the heterogeneity in preferences

¹⁴ The horizontal specification also captures the geography of the urban housing market very naturally, allowing households to have preferences over neighborhoods depending on the distance from their employment locations. This gives rise to variation in the aggregate demand for housing in various neighborhoods throughout the metro area, thereby increasing equilibrium housing prices in neighborhoods near employment centers.

¹⁵ It is important to point out that this flexible feature of our model is made possible because we abstract from issues related to local politics. As Epple, Filimon, and Romer (1993) note, incorporating local politics into models of residential sorting requires restrictions to be placed on preferences in order to guarantee the existence of an equilibrium. Important recent papers by Epple and Sieg (1999) and Epple, Romer and Sieg (2001) estimate equilibrium models that include voting over the level of public goods, restricting households to have shared rankings over a single public goods index. We view our model as having a comparative rather than absolute advantage over the papers in this line of the literature, better suited for an institutional setting such as that which holds in Californian, where Proposition 13 leaves almost no discretion over property tax rates or the level of public goods spending at the local level.

for neighborhood sociodemographic characteristics that will determine the extent to which the initial stratification induced by the variation in school quality across the metropolitan area is reinforced by additional sorting due to preferences for one's neighbors.

Characterizing the Housing Market. As with all models in this literature, the existence of a sorting equilibrium is much easier to establish if the individual residential location decision problem is smoothed in some way. To this end, 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_h^i 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.¹⁶

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

$$(3) \quad V_h^i > V_k^i \quad \Rightarrow \quad W_h^i + \varepsilon_h^i > W_k^i + \varepsilon_k^i \quad \Rightarrow \quad \varepsilon_h^i - \varepsilon_k^i > W_k^i - W_h^i \quad \forall \quad k \neq h$$

¹⁶ For expositional ease and without loss of generality, let the measure of this continuum be one.

where W_h^i includes all of the non-idiosyncratic components of the utility function V_h^i . 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_h^i 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 $\{\mathbf{X}, \mathbf{Z}, \mathbf{p}, \boldsymbol{\xi}\}$:¹⁷

$$(4) \quad P_h^i = f_h(z^i, \mathbf{Z}, \mathbf{X}, \mathbf{p}, \boldsymbol{\xi})$$

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

Aggregating the probabilities in equation (4) over all observed households yields the predicted demand for each housing type h , D_h :

$$(5) \quad D_h = \sum_i P_h^i .$$

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

$$(6) \quad D_h = S_h, \quad \forall h \quad \Rightarrow \quad \sum_i P_h^i = S_h \quad \forall h .$$

¹⁷ 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 (2004a) show that a unique set of prices (up to a scale) clears the market.

When 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.¹⁸ 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 impossible 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. We discuss the issue of uniqueness further in the context of the counterfactual simulations in Section 6 below.

¹⁸ Bayer, McMillan, and Rueben (2004a) establish the existence of a sorting equilibrium as long as (i) the indirect utility function shown in equation (1) is decreasing in housing prices for all households; (ii) indirect utility is a continuous function of neighborhood sociodemographic characteristics; and (iii) ϵ is drawn from a continuous density function.

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 (2003) and included in a technical appendix. It is helpful in describing the estimation procedure to first introduce some notation. In particular, rewrite the indirect utility function as:

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

where

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

and

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

In equation (8), 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 ε_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,¹⁹ returning estimates of the heterogeneous parameters in λ and mean indirect utilities, δ_h . This 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 ε_h^i is drawn from the extreme value distribution, in which case this probability can be written:

$$(10) \quad P_h^i = \frac{\exp(\delta_h + \hat{\lambda}_h^i)}{\sum_k \exp(\delta_k + \hat{\lambda}_k^i)}$$

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

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

¹⁹ Formally, the validity of this first stage procedure requires the assumption that the observed location decisions are individually optimal, given the collective choices made by other households and the vector of market-clearing prices *and* that households are sufficiently small such that they do not interact strategically with respect to particular draws on ε . This ensures that no household's particular idiosyncratic preferences affect the equilibrium and the vector of idiosyncratic preferences ε is uncorrelated with the prices and neighborhood sociodemographic characteristics that arise in any equilibrium. For more discussion, see the Technical Appendix.

where I_h^i 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 $\tilde{\ell}$. Notice that the quasi-likelihood function developed here is based solely on the notion that each household's residential location is optimal given the set of observed prices and the location decisions of other households.

The Endogeneity of School Quality and Neighborhood Sociodemographic Composition. Having estimated the vector of mean indirect utilities in the first stage of the estimation, the second stage of the estimation involves decomposing δ into observable and unobservable components according to the regression equation (8).²⁰ 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 describe the construction of an instrument for price in the Technical Appendix.

A second identification issue concerns the correlation of neighborhood sociodemographic characteristics 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.

There are, however, good reasons to think that households will sort with respect to such boundaries. Thus, while boundary fixed effects are likely to do a good job of controlling for differences in unobserved fixed factors, neighborhood sociodemographics are likely to vary discontinuously at the boundary. In this way, the use of boundary fixed effects isolates 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 district boundary fixed effects when estimating equation (11). 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 district boundary in the metropolitan area (e.g., Palo Alto-Menlo Park).²¹ We show the variation in school quality and neighborhood sociodemographics at school district boundaries in the next section after first describing the basic features of the dataset.

²⁰ Notice that the set of observed residential choices provides no information that distinguishes the components of δ . That is, however δ is broken into components, the effect on the probabilities shown in equation (10) is identical.

²¹ A number of empirical issues arise in incorporating school district boundary fixed effects into our analysis. A central feature of local governance in California helps to eliminate some of the problems that naturally arise with the use of school district boundaries, as Proposition 13 ensures that the vast majority of school districts within California are subject to a uniform effective property tax rate of one percent. Concerning the width of the boundaries, we experimented with a variety of distances and report the results for 0.25 miles, as these were more precise due to the larger sample size.

4 DATA

The analysis conducted in this paper is facilitated by access to restricted Census microdata for 1990. These restricted Census data provide the detailed individual, household, and housing variables found in the public-use version of the Census, but also include information on the location of individual residences and workplaces at a very disaggregate level. In particular, while the public-use data specify the PUMA (a Census region with approximately 100,000 individuals) in which a household lives, the restricted data specify the Census block (a Census region with approximately 100 individuals), thereby identifying the local neighborhood that each individual inhabits and the characteristics of each neighborhood far more accurately than has been previously possible with such a large-scale data set.

We use data from six contiguous counties in the San Francisco Bay Area: Alameda, Contra Costa, Marin, San Mateo, San Francisco, and Santa Clara. We focus on this area for a number of reasons. First, this geographic area is reasonably self-contained, and is sizeable along a number of dimensions, including over 1,100 Census tracts, and almost 39,500 Census blocks, the smallest unit of aggregation in the data. The sample consists of 242,100 households. More importantly, the use of data from California makes it reasonable to use school district rather than school attendance zone boundaries in the analysis. In particular, Proposition 13 ensures that local jurisdictions have almost no discretion over property tax rates or the level of public good spending including school spending. In this way, unlike almost anywhere else in the country, one would not expect much variation in property values across school district boundaries to arise due to

differential property tax rates in California. This same feature the public finance system may also diminish the overall strength of Tiebout-type sorting in California, as households are not free to select different tax rates and local public goods packages in each jurisdiction. For this reason we expect our analysis to generally provide a lower bound on the importance of school-related sorting relative to other states.

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,²² 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. For each of these measures, a detailed description of the process by which the original data were assigned to each house is provided in a Data Appendix. 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.

²² As described in the Data Appendix, we construct a single price vector for all houses, whether rented or owned. Because the implied relationship between house values and current rents depends on expectations about the growth rate of future rents in the market, we estimate a series of hedonic price regressions for each of over 40 sub-regions of the Bay Area housing market. These regressions return an estimate of the

School Assignment and School Quality. While we have an exact assignment of Census blocks to school districts, in the absence of comprehensive information about within-district school attendance areas, we employ the following approach for linking each Census block to a school: For a given Census block, we calculate the distance to each school in its district. We then first assign the Census block to the closest school within its district. Using this closest school assignment, we can then calculate a predicted enrollment for each school (calculated by summing over the school-aged children in each Census block assigned to a school) and compare this measure to the actual enrollment of the school. To correct discrepancies in predicted versus actual enrollment, we then use an intuitive procedure to adjust the assignment of Census blocks to schools so as to ensure that predicted enrollments equal their actual counterparts in each school in each district. We describe this procedure in detail in the Data Appendix; the results are not sensitive to this adjustment.

As our measure of school quality, we use the average test score for each school, averaged over two years. Averaging helps to reduce any year-to-year noise in the measure. When variables that characterize the sociodemographic composition of the school or surrounding neighborhood are included in the analysis, the estimated coefficient on average test score picks up what households are willing to pay for an improvement in average student performance at a school holding the sociodemographic composition constant. While the average test score is an imperfect measure of school quality, it has the advantage of being easily observed by both parents and researchers and

ratio of house values to rents for each of these sub-regions and we use the average of these ratios for the Bay Area, 264.1, to convert monthly rent to house value for the purposes of reporting results at the mean.

consequently has been used in most analyses that attempt to measure demand for school quality.

Boundaries. 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 0.25 miles of a school district boundary.²³ Comparing the first column to the third column of the table, it is immediately obvious that the houses near school district boundaries are not fully representative of those in the Bay Area as a whole. To address this problem, we create sample weights for the houses near the boundary.²⁴ Column 7 of Table 1 shows the resulting weighted means, indicating that using these weights makes the sample near the boundary much more representative of the full sample.

The fourth and fifth columns report means for houses within 0.25 miles of a boundary, comparing houses on the high versus low average test score side of the each boundary; the sixth column reports t-tests for the difference in means. Comparing these differences reveals that houses on the high side cost \$53 more per month and are assigned to schools with test scores that are 43-point higher on average.²⁵ 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-

²³ We experimented with a variety of distances and report the results for 0.25 miles, as these were more precise due to the larger sample size.

²⁴ The following procedure is used: we first regress a dummy variable indicating whether a house is in a boundary region on the vector of housing and neighborhood attributes using a logistic regression. Fitted values from this regression provide an estimate of the likelihood that a house is in the boundary region given its attributes. We use the inverse of this fitted value as a sample weight in subsequent regression analysis conducted on the sample of houses near the boundary.

²⁵ As described in the Data Appendix, we construct a single monthly price vector for all houses, whether rented or owned.

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 control 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.²⁶

5 PARAMETER ESTIMATES

Estimation of the full model proceeds in two stages. The first stage recovers interaction parameters and a vector of mean indirect utilities; the second stage returns the components of mean indirect utility. To give the reader a sense of the interplay between the estimates of the demand for school quality and sociodemographic characteristics of neighbors, we report results for two main specifications, which respectively exclude and include variables that characterize the racial composition, average educational attainment, and average income of the neighborhood (Census block group). To make the discussion of these estimates more transparent, we transform the estimates so that they can be described in terms of marginal willingness-to-pay measures (MWTP), reporting these estimates in Tables 2-4.

Table 2 reports the implied measures of the mean MWTP for school average test scores and other selected housing and neighborhood attributes for six specifications of the

²⁶ In terms of the estimates related to neighborhood sociodemographic characteristics, the key point about using school district boundary fixed effects rather than Census tract fixed effects is that in the boundary case we have a clear sense of what fundamentally leads to the sorting of households across neighborhoods within the region upon which the fixed effect is based. Because we control directly for that cause of the sorting - schooling in this case - we are less concerned that the variation in sorting is related to variation in unobservables within the region upon which the fixed effect is based.

mean indirect utility regression.²⁷ Results are reported for the two main specifications, which exclude and include neighborhood sociodemographic variables, respectively. For each of these main specifications, results are reported for the full sample and for a sample of houses within 0.25 miles of school district boundaries, with and without including boundary fixed effects. In all cases, when the sample of houses is restricted to those within 0.25 miles of a boundary, sample weights (as described in Section 4) are used in order to make this sample as close to representative of the full sample as possible. Comparing columns 1 and 2 or columns 4 and 5, it is clear that results are very similar whether the full sample or the weighted sub-sample of houses near a school district boundary is used.

The first three columns of Table 2 report results for specifications that exclude neighborhood sociodemographics. In this case, the estimated mean MWTP for a one standard deviation increase in school average test score declines from \$123 to \$82 in monthly rent (\$21,500 in house value) when boundary fixed effects are included in the analysis. The final three columns report results for analogous specifications that include neighborhood sociodemographic characteristics. The coefficient on the average test score in these specifications returns the average of what households are willing to pay for a standard deviation increase in the average test score *conditional* on the sociodemographic characteristics of the neighborhood, which are in most cases also indicative of the sociodemographic characteristics of the local school. The estimated MWTP for a one standard deviation increase in the average test score in this case

²⁷ The specifications of the mean indirect utility regressions are reported in Appendix Table 1. The mean MWTP measures reported in Table 2 are calculated by dividing the coefficient associated with each choice characteristic by the coefficient on price.

declines to \$26 in monthly rent (\$6,900 in house value), which is approximately 2.4 percent of the average house price index for our Bay Area sample.²⁸

These results make clear that much of what initially appears in the specification without neighborhood sociodemographics to be a significant willingness on the part of households to pay for school quality is instead more properly attributed to the characteristics of neighbors or peers. That the resulting MWTP for school average test scores is relatively small is exactly what one would expect if households have difficulty inferring the quality of a school from published average test score data.²⁹ That is, one would expect households to place a relatively small weight on this measure when choosing neighborhoods if the signal that the published average test score provided about actual school quality were small relative to the noise that it contains related to differences in the underlying composition of individuals taking the test. In fact, some of the weight that parents place directly on neighborhood sociodemographics may result from a belief that these provide a better indication than does the test score itself of the quality of the education that their children will receive in the local schools, especially if parents perceive peer effects to be important.

Before turning to the results related to heterogeneity in preferences, it is important to point out that the final two columns of Table 2 also show the impact of including boundary fixed effects on the estimates of mean preferences for neighborhood sociodemographic characteristics. Comparing these columns reveals the pattern of results that one would expect if boundary fixed effects control in part for unobserved

²⁸ This is very similar to the estimates of Black (1999) and Kane, Staiger, and Riegg (2004).

²⁹ This is especially true in 1990, which pre-dates most concerted efforts on the part of states to provide information to households about the quality of the local school. In many cases, such information simply reflects that average test score anyway.

neighborhood quality and unobserved quality is positively correlated with neighborhood income and education and negatively correlated with the fraction of non-white households.³⁰ Thus boundary fixed effects seem to be effective in controlling for fixed aspects of unobserved neighborhood quality that are correlated with neighborhood sociodemographics, and thus provide an attractive way of estimating preferences for both school quality *and* neighborhood sociodemographic characteristics in the presence of this important endogeneity problem.

Heterogeneity in Willingness-to-Pay

The interaction parameters estimated in the first stage for the specifications that exclude and include neighborhood sociodemographics, respectively, are reported in Appendix Tables 2 and 3. These specifications correspond to the mean MWTP estimates reported in columns 3 and 6 of Table 2, respectively. In each case, the model controls simultaneously for the effect of each of a series of household characteristics (income, education, race, work status, age, and household structure) on the marginal willingness-to-pay for a variety of housing and neighborhood attributes, including the average test score of the local school. The model also captures the spatial aspect of the housing market by allowing households to have preferences over commuting distance.³¹

³⁰ The fact that the estimated coefficient on the average test score rises from \$20 to \$26 when boundary fixed effects are included reflects that fact it is positively correlated with neighborhood income and education and negatively correlated with the fraction of non-white households. Thus, the estimated coefficient on the average test score tends to rise as the coefficients on these other variables change, as they do in moving from column 5 to column 6 in Table 2.

³¹ We treat a household's primary work location as exogenous, calculating the distance from this location to the location of the neighborhood in question. MWTP estimates for other housing and neighborhood attributes based on a specification without commuting distance are qualitatively similar except for variables that are strongly correlated with employment access, such as population density.

Table 3 converts the estimates of the heterogeneity in MWTP for the average test score into dollar terms. The two columns of the table report results for specifications that exclude and include neighborhood sociodemographic characteristics, respectively. The first row of this table reports the estimated mean MWTP for the change listed in each column heading: for example, the first entry reports the \$82 mean MWTP for a one standard deviation increase in the average test score conditional on neighborhood sociodemographics initially reported in column 3 of Table 2. The remaining rows report the difference in MWTP associated with the comparison of household characteristics shown in the row heading. Thus, the second row shows how a household's MWTP changes with an increase of \$100,000 in income.

Analogously to the effect of including neighborhood sociodemographic variables on the estimates of the mean MWTP reported in Table 2, including these variables decreases the magnitude of the estimated heterogeneity in demand for school quality, decreasing the coefficient associated with household income by 60 percent, with education by 75 percent, and with race (which may proxy in part for wealth) by upwards of 80 percent. The estimates of the specification that includes neighborhood sociodemographics also returns the expected positive additional MWTP associated with the presence of school-aged children. Table 4 reports analogous measures of the heterogeneity in MWTP for selected housing and neighborhood attributes from the specification that includes neighborhood sociodemographic characteristics, revealing for example that increases in household income are typically associated with large increases in the demand for housing attributes, and that college educated households are willing to pay a substantial premium relative to less educated households to live with more

educated neighbors. Specifically, they are willing to pay around \$59 per month more than less educated households for a 10 percent increase in college-educated neighbors. Not surprisingly, strong racial interactions are also at work in the housing market, leading to significant sorting along this dimension as well.³²

The pattern of heterogeneity revealed in Tables 3 and 4 is exactly the type of preference structure that gives rise to a substantial social multiplier in considering the effect of an exogenous school quality on neighborhood stratification. In particular, given the heterogeneity in preferences for school quality, Asian and white households and households with more education and income are generally more likely to sort into a neighborhood following an increase in school quality there. This initial effect on neighborhood composition then gives rise to additional re-sorting according to the preferences directly associated with neighborhood sociodemographic characteristics. This additional re-sorting works to reinforce the initial changes, multiplying the effects of the exogenous change in school quality up by a factor that may be sizeable given the magnitude of the parameters associated with neighborhood sociodemographics.

6 SCHOOL-RELATED SORTING AND NEIGHBORHOOD STRATIFICATION

To explore the impact of school-related sorting on neighborhood stratification, we now turn to the main economic analysis of the paper. In particular, given the set of preference parameters estimated above, we simulate the housing market equilibrium in a

³² Note that the strength of the estimated racial interactions reported in Table 4 may reflect the presence of centralized discriminatory practices in the market in addition to the direct preferences of households to live with others of the same race. In this way, the counterfactual that we conduct below should be interpreted as

counterfactual world in which preferences for the average test score are set identically equal to zero for all households. The resulting reductions in stratification of households across neighborhoods on the basis of education and income provide measures of the general equilibrium impact of school-related sorting on stratification. Before describing these simulations in detail, however, we begin this section by discussing the pattern of neighborhood stratification in the Bay Area as well as the corresponding measures predicted by our model, providing a gauge of how well the model fits the data in this dimension.

Initial Measures of Neighborhood Stratification. To quantify the extent of neighborhood stratification in the Bay Area, we use a series of simple measures based on the average sociodemographic composition of the neighborhoods in which households with given characteristics live.³³ Using the underlying average exposure of households in each education category to one another, for example, we measure the level of neighborhood stratification by calculating the degree to which households with a high school degree or less and a college degree or more, respectively, are ‘over-exposed’ (i.e., exposed to a fraction greater than the average for the Bay Area as a whole) to households with the same education level.

To see these calculations more explicitly, consider the top panel (Panel A) of Table 5, which constructs two measures of neighborhood stratification on the basis of education using the sample of nearly a quarter of a million households. Households are

holding the structure of racial interactions constant, regardless of the source. See Bayer, McMillan, and Rueben (2004b) for more discussion of this issue.

³³ We define neighborhoods on the basis of (synthetic) school attendance zones, as described in Section 4. Defining neighborhoods on the basis of alternative neighborhood definitions - Census block groups for

first assigned to one of three mutually exclusive education categories based on the educational attainment of the household head - high school diploma or less, some college, college degree or more.³⁴ These individual education categories are given in the first three column headings of the panel, while the fourth column considers all households. Each column of Panel A gives the respective average neighborhood education distribution for each type of household. The figures in the first column, for example, imply that households with a high school degree or less live on average in neighborhoods in which 41.2 percent of household heads (not counting the original household) have a high school degree or less, 22.6 have had some college, and 36.2 have a college degree or more. The fifth column of Panel A reports the extent to which households are ‘over-exposed’ to households in the same education category. Thus, in the Bay Area, households with a HS degree or less are exposed on average to 21.8 percent more households with a HS degree or less than in the population as a whole. Similarly, households with a college degree or more are exposed to 14.7 percent more households with a college degree than in the population as a whole

Panel A of Table 6 reports analogous measures related to income stratification. In this case, households are assigned to categories on the basis of income quintile. To match the education categories as much as possible, we form three categories: the bottom two quintiles, the middle quintile, and the top two quintiles. For our sample of 242,100 households, households in the bottom two income quintiles are exposed on average to 17.4 percent more households in the bottom 40 percent of the income distribution than

example - tends to give rise to more marked residential stratification, primarily because block groups are smaller geographic units.

³⁴ In what follows, when we refer to households as having a college degree, this should be taken to mean that the household head as identified in the Census has a college degree, and similarly for other categories.

are in the Bay Area as a whole, while the analogous figure is 18.1 percent for households in the top 40 percent of the income distribution.

To provide a sense of how well the estimated model fits the data, Panel B of Tables 5 and 6 reports the same pattern of average exposure rates based on the predictions of the model.³⁵ The model used to calculate the measures reported in Tables 5 and 6 is the full specification including neighborhood sociodemographics, which corresponds to the MWTP estimates reported in column 6 of Table 2, column 2 of Table 3, and Table 4. Panel B has the same basic structure as the top panel, and the comparison makes clear that our model predictions are reasonably close to the sample measures; on average, the nine predicted exposure measures reported in each table are slightly more than half a percentage point off from the sample measures. Thus our model does a very reasonable job of matching stratification patterns in the data, even though the estimation procedure does not fit neighborhood compositions directly. These pre-simulation predictions of the model play an important role in the analysis that follows, providing the appropriate benchmark against which the equivalent measures generated in the simulations should be compared.

Counterfactual Simulations – Computation and Basic Setup. The idea behind the simulations that we conduct is intuitive: having estimated household preferences for the average test score, along with other choice characteristics, we suppose (counterfactually) that households no longer care about school average test scores when deciding where to

³⁵ Specifically, the model provides predictions of the probability that each household occupies each housing type in its choice set in equilibrium. These probabilities are aggregated up to the neighborhood level to provide predicted neighborhood compositions, which can then be averaged by household education and income for comparison with the calculations reported in Panel A.

live. We then recalculate neighborhood compositions, generating new stratification measures that correspond to those reported in Panels A and B of Tables 5 and 6.

The basic structure of the computation of the new equilibrium consists of a loop within a loop. Having set the preference parameters associated with the average test score to zero, we first calculate a new set of prices that clears the market; Berry (1994) ensures that there is a unique set of market clearing prices up to scale. Using these new prices and the initial sociodemographic composition of each neighborhood, we calculate the probability that each household chooses each housing type, and aggregating these choices to the neighborhood level, the corresponding predicted sociodemographic composition of each neighborhood. We replace the initial neighborhood sociodemographic measures with these new measures and start the loop again – i.e., calculate a new set of market clearing prices with these updated neighborhood sociodemographic measures. We continue this process until the neighborhood sociodemographic measures converge. The household location decisions corresponding to the final sociodemographic measures along with the vector of housing prices that clears the market then represents the new equilibrium.³⁶

Because our model is fundamentally static, this equilibrium should be viewed as representing the way the housing market equilibrium in the Bay Area would look in a

³⁶ As discussed in Section 2, uniqueness is not a generic property of our sorting model. Without this property, it is sometimes difficult to justify counterfactual simulations corresponding to non-marginal changes in the primitives of the underlying model. As argued by Debreu (1969), however, the property of local uniqueness provides a coherent basis for conducting counterfactual simulations associated with a *marginal* change in the model's primitives. In this case, the results of the type of equilibrium counterfactual simulations used here correspond to a series of GE comparative static measures estimated at the current equilibrium. To alleviate concerns about the role of multiple equilibria in our analysis, we repeated the calculations for a marginal change in the model's primitives (a 5 percent reduction in the magnitude of preference parameters associated with average test score). The results were qualitatively similar to those reported here; we choose to report the results of a non-marginal change corresponding to setting the test score preference parameters to zero because it corresponds more directly to the thought experiment that we would like to carry out.

world in which school quality never mattered, rather than as the outcome if school quality was suddenly eliminated as a factor in residential location decisions following some kind of new policy.³⁷ Accordingly, the difference between the level of neighborhood stratification in the current state of the world and this new equilibrium represents a measure of the *general equilibrium* effect of school-related sorting on neighborhood stratification. In addition to this general equilibrium effect, we also report a *partial equilibrium* effect that measures the initial predicted change in stratification following the elimination in preferences related to the average test score. The partial equilibrium effects reported below do not account for the reinforcing effects on stratification associated with the additional re-sorting of households on the basis of sociodemographic characteristics.

Counterfactual Simulations – Results. Having set the preference parameters associated with average test score to zero, Panels C and D of Tables 5 and 6 present the predicted pattern of neighborhood stratification under this counterfactual, reporting partial and general equilibrium measures respectively. Panel C shows that the initial effect of setting the preference parameters associated with the average test score to zero is an immediate decline in the rate of over-exposure of households with a high school degree or less to each other, falling from 23.7 percent to 21.2 percent. As the sixth column of Table 5 reports, this amounts to a 10.9 percent reduction in over-exposure. Similarly, households with a college degree initially experience about a 6 percent decline in the rate at which

³⁷ In general, incorporating dynamic behavior into an estimable model that allows for sorting on the basis of neighborhood sociodemographics is very difficult; such a model has yet to be developed in the literature. Also, our general equilibrium predictions are fully based on Tiebout assumptions - see Ferreira (2004) for incorporating moving costs into the estimation of MWTP for housing amenities.

they are over-exposed to one another. Panel C of Table 6 reports a similar set of partial equilibrium results for income stratification, with the rate of over-exposure of households in the bottom and top 40 percent of income distribution respectively to households in the same part of income distribution falling by 10-12 percent.

These partial equilibrium predictions capture the direct effect of eliminating preferences for the average test score, holding constant the initial sociodemographic composition of each neighborhood. Thus, in the current state of the world, these measures provide a clear sense of how much less stratified the metropolitan area would be if direct preferences over school average test scores (holding neighborhood composition constant) were eliminated. Given the magnitude of the heterogeneous preferences for the characteristics of one's neighbors estimated above, however, the resulting changes in neighborhood sociodemographic compositions will in turn be reinforced by additional re-sorting in response to these sociodemographic changes. In particular, in the subsequently less stratified world, differences in preferences for neighborhood characteristics simply have a smaller role in the location decision of each household, thereby leading to even less stratification.

Starting with the initial change in neighborhood compositions following the elimination of preferences for school average test scores, Panel D of Tables 5 and 6 reports the levels of neighborhood stratification once a new equilibrium has been reached. As the corresponding stratification measures make clear, the full general equilibrium effect of eliminating preferences for school quality is substantially greater than the partial equilibrium effect. The rate of over-exposure of households with a high school degree or less to one another, for example, falls from 23.7 percent to 14.5 percent

- a 38.8 percent decline. Similarly, the rate of over-exposure of college educated households to one another falls by over 25 percent. The comparable numbers for income stratification also average over 25 percent.³⁸

In sum, the general equilibrium measures of the reduction in stratification following the elimination of preferences for school average test scores are almost *four* times greater than the comparable partial equilibrium measures for education stratification and more than *two* times greater for income stratification. The especially strong general equilibrium effect for education stratification can be linked to the underlying coefficient estimates. As Table 4 indicates, households with a college degree are willing to pay substantially more than those without a college degree to live with college-educated neighbors. Thus the initial decline in neighborhood stratification induced by the elimination of preferences for test scores has an especially strong subsequent effect on residential sorting on the basis of education.

As mentioned in Section 5, it is important to point out that some of the weight placed directly on neighborhood sociodemographic characteristics as opposed to average test scores in the utility function may be related to inference issues related to school quality. That is, parents may perceive sociodemographic characteristics to provide a better indication of school quality than the average test score itself. In this case, in conducting the counterfactual simulations above, one would want to simultaneously eliminate the preference parameters that multiply school average test scores and reduce the magnitude of the parameters associated with neighborhood sociodemographic characteristics. This would serve to increase both the partial and general equilibrium

³⁸ The results of this analysis are not sensitive to the measures of stratification used here. Qualitatively similar reductions in stratification, for example, were obtained using less aggregate groupings of

estimates of the portion of stratification attributable to school-related sorting. While the difference between the general and partial equilibrium results would be smaller in this case, the central point that we make in this paper would be strengthened: the full impact of school-related sorting on neighborhood stratification is substantially greater than a narrow examination of the direct estimates of the demand for school average test scores would suggest.

7 CONCLUSION

In recent years, the empirical literature in the economics of education has reached a consensus that the direct effect of school quality improvements on household demand is relatively small, while the demand for peers and neighbors explains a great deal of observed sorting across schools and neighborhoods. This paper considers the economic content of this consensus, demonstrating that the full effect of school-related sorting on neighborhood stratification is in fact substantially greater than direct estimates of demand for school quality would seemingly imply. This much greater general equilibrium effect comes about precisely because certain households place such an important weight on the characteristics of their neighbors when choosing where to reside and where to send their children to school. Thus sorting directly related to neighborhood sociodemographic characteristics works to reinforce any stratification initially induced by exogenous variation in school quality.

To quantify the magnitude of this general equilibrium effect, we estimate a flexible model of residential sorting using restricted Census data that provide the precise residential locations of nearly a quarter of million households. These data permit

individuals by income and education. Those calculations are available from the authors upon request.

flexibility both in the model itself (we estimate a horizontal model of sorting that includes many housing and neighborhood attributes and allows the demand for each of these attributes to vary with a wide set of household characteristics) and in the way that we use the data to estimate it. Here, we incorporate the boundary fixed effects approach pioneered by Black (1999) in estimating the parameters of our broader sorting model. The estimated model becomes a powerful tool for conducting general equilibrium counterfactual simulations, providing an intuitive way to estimate the full impact of school-related sorting on neighborhood stratification.

Consistent with the recent empirical literature, we estimate that households are willing to pay an average of only about 2 percent in house price for a one standard deviation increase in school average test scores. At the same time, there is heterogeneity in willingness to pay for improvements in quality; and further, households have strongly heterogeneous preferences for the characteristics of their neighbors; college-educated households, for example, are willing to pay substantially more than others for an increase in the fraction of college-educated households in their neighborhood. Using the estimated model, we then calculate measures of neighborhood stratification in a counterfactual environment in which household preferences related to school average test scores are set identically equal to zero. The results of this exercise reveal that school-related sorting explains roughly 25 percent of the neighborhood stratification on the basis of income and 30 percent on the basis of education. These full general equilibrium effects are on average four times greater than corresponding partial equilibrium measures in the case of education stratification and twice as great in the case of income. Partial equilibrium measures consider only the direct effect of eliminating preferences related to

school quality - that is, do not allow for additional re-sorting in response to the resulting changes in neighborhood sociodemographic compositions.

Our findings indicate that the consensus reached in the empirical literature regarding preferences for schools and neighbors is considerably more supportive of the long theoretical tradition in local public finance related to Tiebout-type sorting than the direct estimates of the demand for school average test scores might suggest. School-related sorting has quantitatively significant effects on neighborhood stratification once important general equilibrium considerations are taken into account. Moreover, these results also imply that the effect of Tiebout sorting on the competitiveness of the education marketplace is substantially greater than direct estimates of the elasticity of demand with respect to an exogenous change in school quality would imply. Again in this case, the direct effect of a change in school quality on demand would be magnified by the effect of this change on who sorts into the corresponding school and neighborhood.

REFERENCES

Bayer, Patrick, Fernando Ferreira, and Robert McMillan (2003), "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," Economic Growth Center, Yale University Working Paper No. 872.

Bayer, Patrick, Robert McMillan, and Kim Rueben, (2004a), "An Equilibrium Model of an Urban Housing Market," unpublished manuscript, Yale University.

_____, (2004b), "Residential Segregation in General Equilibrium," Economic Growth Center, Yale University Working Paper No. 885.

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*, Vol 63, pp. 841-890.

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

Cullen, Julie, Brian Jacob, and Steven Levitt, (2003), "The Effect of School Choice on Student Outcomes: Evidence from Randomized Lotteries," NBER Working Paper No. 10113.

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 Debreau*, Cambridge University Press, 1983.

Epple, Dennis, David Figlio, and Richard Romano, (2004), "Competition between Private and Public Schools: Testing Stratification and Pricing Predictions," *Journal of Public Economics* 88, 1215-45.

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*, Vol. 24, pp. 281-304.

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

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*, Vol. 69, No. 6., pp. 1437-1455.

Epple, Dennis and Holger Sieg, (1999), "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy*, Vol. 107, No. 4., pp. 645-681.

Fernandez, Raquel and Richard Rogerson, (1996), "Income Distribution, Communities, and the Quality of Public Education." *Quarterly Journal of Economics*, Vol. 111, No. 1., pp. 135-164.

_____, (2003), "Equity and Resources: An Analysis of Educational Finance Systems," *Journal of Political Economy*, vol. 111, no. 4.

Ferreira, Fernando V. (2004), "*You Can Take It with You*: Transferability of Proposition 13 Tax Benefits, Residential Mobility, and Willingness to Pay for Housing Amenities." Center for Labor Economics, University of California, Berkeley, Working Paper No. 72.

Ferreira, Maria, (2003), "Estimating the Effects of Private School Vouchers in Multi-District Economies," mimeo, Carnegie-Mellon University.

Hoxby, Caroline M., (2000) "Does Competition Among Public Schools Benefit Students and Taxpayers?" *American Economic Review*, 90 (5), 1209-38.

Kane, Thomas, Douglas Staiger, and Stephanie Riegg, (2004) "Do Good Schools or Good Neighbors Raise Property Values?" mimeo, University of California, Los Angeles.

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.

Rothstein, Jesse, (2003), "Good Principals or Good Peers? Parental Valuation of School Characteristics, Tiebout Equilibrium, and the Incentive Effects of Competition Among Jurisdictions," mimeo, Princeton University.

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

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

Timmins, Christopher (2003), "If You Can't Take the Heat, Get Out of the Cerrado... Recovering the Amenity Cost of Non-Marginal Climate Change in Brazil," mimeo, Yale University.

Walsh, Randall (2004), "Endogenous Open Space Amenities in a Locational Equilibrium," CEA Working Paper No. 04-03, February.

Table 1. Full Sample, and School District Boundary Sub-Sample

Sample Boundary/Weights Observations	full sample		within 0.25 miles of boundaries				
			actual sample	high test score side*	low test score side*	t-test for	weighted sample
	242,100		27,958	13,348	14,610	difference in	27,958
	(1)	(2)	(3)	(4)	(5)	means	(6)
	Mean	S.D.	Mean	Mean	Mean	((4) versus (5))	Mean
<u>Housing/Neighborhood Characteristics</u>							
monthly house price	1,087	755	1,130	1,158	1,105	5.71	1,098
average test score	527	74	536	558	515	50.96	529
1 if unit owned	0.597	0.491	0.629	0.632	0.626	1.04	0.616
number of rooms	5.114	1.992	5.170	5.207	5.134	3.13	5.180
1 if built in 1980s	0.143	0.350	0.108	0.118	0.099	5.09	0.148
1 if built in 1960s or 1970s	0.391	0.488	0.424	0.412	0.437	4.22	0.406
elevation	210	179	193	194	192	1.14	212
population density	0.434	0.497	0.352	0.349	0.355	2.08	0.374
crime index	8.184	10.777	6.100	6.000	6.192	2.36	7.000
% Census block group white	0.681	0.232	0.704	0.712	0.686	9.62	0.676
% Census block group black	0.081	0.159	0.071	0.065	0.076	6.21	0.080
% Census block group Hispanic	0.110	0.114	0.113	0.107	0.119	8.62	0.117
% Census block group Asian	0.122	0.120	0.112	0.110	0.113	2.50	0.121
% block group college degree or more	0.438	0.196	0.457	0.463	0.451	5.14	0.433
average block group income	54,744	26,075	57,039	58,771	55,457	10.23	55,262
<u>Household Characteristics</u>							
household income	54,103	50,719	56,663	58,041	55,405	4.20	55,498
1 if children under 18 in household	0.333	0.471	0.324	0.322	0.325	0.54	0.336
1 if black	0.076	0.264	0.066	0.062	0.070	2.69	0.076
1 if Hispanic	0.109	0.312	0.111	0.102	0.119	4.54	0.115
1 if Asian	0.124	0.329	0.112	0.114	0.110	1.06	0.121
1 if white	0.686	0.464	0.706	0.717	0.696	3.86	0.682
1 if less than high school	0.154	0.361	0.141	0.134	0.147	3.12	0.152
1 if high school	0.184	0.388	0.176	0.177	0.175	0.44	0.183
1 if some college	0.223	0.417	0.222	0.222	0.223	0.20	0.225
1 if college degree	0.291	0.454	0.294	0.295	0.294	0.18	0.286
1 if more than college	0.147	0.354	0.166	0.172	0.161	2.46	0.155
age (years)	47.607	16.619	47.890	48.104	47.699	1.99	47.660
1 if working	0.698	0.459	0.705	0.702	0.709	1.28	0.701
distance to work (miles)	8.843	8.597	8.450	8.412	8.492	0.82	8.490

Notes: Columns 1 and 2 report the mean and standard deviation for key variables for the full sample. Column 3 reports means for the sample of houses within 0.25 miles of a school district boundary. Columns 4 and 5 report means on the high versus low test score side of boundaries. Column 6 provides a t-statistic for a test of whether the means reported in columns 4 and 5 are equal. Column 7 reports weighted means for the sample of houses within 0.25 miles of a school district boundary. Weights are constructed so as to make the boundary sample more representative of the full sample and are described in the main text. In constructing columns 4 and 5, we assign each house in the full sample to the nearest school district boundary, noting whether its local school has a higher test score than the school associated with the closest Census block on the other side of the boundary.

Table 2: Implied Mean MWTP Measures

Sample	Without Neighborhood Sociodemographics			With Neighborhood Sociodemographics		
	full sample	within .25 mile of boundaries		full sample	within .25 mile of boundaries	
	No	No	Yes	No	No	Yes
Boundary Fixed Effects						
Observations	242,100	27,958	27,958	242,100	27,958	27,958
	(1)	(2)	(3)	(4)	(5)	(6)
average test score (in standard deviations)	126.08	122.89	81.53	20.17	20.19	26.22
	(1.96)	(5.36)	(7.72)	(1.72)	(4.77)	(6.13)
1 if unit owned	209.76	178.37	184.54	165.38	150.77	161.05
	(3.29)	(8.99)	(11.39)	(3.19)	(8.76)	(9.24)
number of rooms	148.98	149.36	138.71	122.03	121.12	118.93
	(1.51)	(4.24)	(5.49)	(1.48)	(4.23)	(4.40)
1 if built in 1980s	129.93	74.74	106.17	99.69	85.50	95.55
	(3.94)	(10.87)	(14.41)	(3.79)	(10.69)	(11.84)
1 if built in 1960s or 1970s	28.48	9.46	15.39	13.79	7.40	4.50
	(2.78)	(8.03)	(10.48)	(2.67)	(7.71)	(8.51)
elevation (/100)	21.09	-4.82	46.46	-1.06	-18.04	12.83
	(0.81)	(2.48)	(6.35)	(0.75)	(2.46)	(5.04)
population density	-100.43	-153.53	-133.08	19.41	41.68	30.33
	(4.23)	(15.64)	(23.85)	(4.30)	(15.76)	(20.09)
crime index	-2.95	-2.30	1.78	0.00	-1.39	1.96
	(0.18)	(0.70)	(2.20)	(0.20)	(0.81)	(1.91)
% Census block group black				-324.67	-318.83	-267.08
				(10.14)	(32.15)	(39.84)
% Census block group Hispanic				-4.42	18.06	138.95
				(14.35)	(46.87)	(63.13)
% Census block group Asian				-97.39	-96.22	155.27
				(11.15)	(37.39)	(55.73)
% block group college degree or more				286.02	206.02	137.71
				(10.50)	(30.58)	(44.53)
average block group income				87.08	96.11	87.61
				(1.25)	(3.86)	(4.00)
F-statistic for boundary fixed effects			5.349			4.162

Note: Specifications shown in the table also include controls for land use (% industrial, % residential, % commercial, % open space, % other) in 1, 2, 3, 4, and 5 mile rings around location and six variables that characterize the housing stock in each of these rings. MWTP measures are reported in terms of a monthly house price. Standard errors are in parentheses.

Table 3. Heterogeneity in Marginal Willingness to Pay for School Average Test Score

	Without Neighborhood Sociodemographics in Model	With Neighborhood Sociodemographics in Model
	<i>One Standard Deviation Increase in Average Test Score</i>	
Mean MWTP	81.53 (7.72)	26.22 (6.13)
Heterogeneity in MWTP		
Household Income (+\$100,000)	40.45 (0.28)	15.66 (0.35)
Children Under 18 vs. No Children	-11.90 (3.07)	7.10 (3.78)
Black vs. White	-93.84 (5.15)	-18.05 (7.50)
Hispanic vs. White	-40.75 (4.64)	-4.64 (5.80)
Asian vs. White	-9.08 (3.99)	5.79 (5.08)
College Degree or More vs. Some College or Less	57.65 (3.46)	14.12 (4.24)
Householder Working vs. Not Working	1.91 (3.16)	6.63 (4.02)
Age (+10 years)	1.02 (0.09)	0.86 (0.11)

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. Standard errors are in parentheses.

Table 4. Heterogeneity in Marginal Willingness to Pay for Select Housing and Neighborhood Attributes*Specification Includes Neighborhood Sociodemographic Characteristics*

	House Characteristics			Neighborhood Sociodemographics				
	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
Mean MWTP	161.05 (9.24)	118.93 (4.40)	95.55 (11.84)	-26.71 (3.98)	13.90 (6.31)	15.53 (5.57)	13.77 (4.45)	87.61 (4.00)
Household Income (+\$100,000)	218.37 (7.13)	61.19 (1.70)	105.07 (7.62)	-15.32 (3.89)	7.73 (4.13)	-0.50 (2.54)	26.17 (2.18)	15.44 (1.13)
Children Under 18 vs. No Children	-12.87 (6.67)	40.06 (1.80)	-24.52 (7.94)	10.38 (2.70)	15.03 (3.41)	12.17 (2.51)	-14.18 (2.15)	5.05 (1.06)
Black vs. White	-63.55 (13.25)	1.56 (3.40)	2.95 (16.38)	96.82 (3.62)	46.13 (5.75)	48.02 (4.84)	16.99 (4.40)	-0.45 (2.27)
Hispanic vs. White	-6.44 (9.53)	-14.14 (2.63)	-8.07 (12.00)	28.89 (3.68)	81.36 (4.01)	18.01 (3.81)	5.43 (3.19)	2.07 (1.41)
Asian vs. White	113.65 (8.96)	-32.92 (2.27)	43.94 (10.77)	27.74 (3.64)	21.95 (4.67)	92.49 (2.78)	-0.05 (2.91)	1.99 (1.41)
College Degree or More vs. Some College or Less	33.83 (7.67)	4.50 (2.05)	42.06 (9.57)	8.34 (3.27)	-4.16 (3.94)	-12.70 (2.91)	59.29 (2.37)	3.66 (1.29)
Householder Working vs. Not Working	42.72 (7.31)	3.69 (1.94)	60.60 (8.92)	-4.71 (2.88)	-2.17 (3.65)	-2.81 (2.82)	-12.62 (2.27)	3.88 (1.04)
Age (+10 years)	6.49 (0.21)	0.30 (0.06)	-2.07 (0.25)	-0.15 (0.08)	-0.56 (0.10)	-0.03 (0.08)	-0.12 (0.06)	0.11 (0.03)

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. Standard errors are in parentheses.

Table 5: Education Stratification Measures Based on Full Model With Neighborhood Sociodemographic:

Panel A: Sample					<u>Over-Exposure to Same Education Level</u>	<u>Percent Reduction in Over-Exposure</u>
<i>Household Characteristic</i>						
<i>Neighborhood Education Composition</i>	HS or less	Some Coll	Coll Deg or more	Overall		
high school or less	41.2	34.2	28.0	33.8	21.8%	
some college	22.6	23.1	21.8	22.3		
college degree or more	36.2	42.7	50.3	43.8	14.7%	
Panel B: Pre-Simulation (Model's Predictions)						
<i>Household Characteristic</i>						
<i>Neighborhood Education Composition</i>	HS or less	Some Coll	Coll Deg or more	Overall		
high school or less	41.9	34.8	27.2	33.8	23.7%	
some college	23.0	23.2	21.4	22.3		
college degree or more	35.2	42.0	51.4	43.8	17.4%	
Panel C: Post-Simulation, Partial Equilibrium						
<i>Household Characteristic</i>						
<i>Neighborhood Education Composition</i>	HS or less	Some Coll	Coll Deg or more	Overall		
high school or less	41.0	35.0	27.7	33.8	21.1%	10.9%
some college	23.1	23.2	21.3	22.3		
college degree or more	35.9	41.8	51.0	43.8	16.3%	6.1%
Panel D: Post-Simulation, General Equilibrium						
<i>Household Characteristic</i>						
<i>Neighborhood Education Composition</i>	HS or less	Some Coll	Coll Deg or more	Overall		
high school or less	38.7	35.1	29.4	33.8	14.5%	38.8%
some college	23.2	23.4	21.1	22.3		
college degree or more	38.1	41.5	49.5	43.8	12.9%	25.9%

Notes (Columns 1-4): The first four columns of each panel of the table report the average neighborhood education distribution for households with the characteristics shown in the column heading. The first column, for example, shows the average education distribution in the neighborhoods in which households with a high school degree or less reside. Panel A reports these measures for the sample; Panel B reports predicted measures using the estimated parameters of the model; Panels C and D report measures based on a counterfactual simulation that sets preferences for the average test score identically to zero for all households. Panel C gives partial equilibrium measures that account for only the initial re-sorting; Panel D reports general equilibrium measures, which are derived from a new equilibrium in which households are allowed to re-sort on the basis of the new neighborhood sociodemographic compositions that follow the initial re-sorting induced by the elimination of preferences for school quality.

Notes (Column 5-6): The fifth column of the table reports the extent to which households with a high school degree or less and a college degree or more are 'over-exposed' to other households with the same education level, respectively. So, for example, the first entry of this column is formed as $(41.2-33.8)/33.8$. The final column shows the percentage reduction in these measures in the simulation relative to the model's predictions, thus providing measures of the portion of neighborhood stratification on the basis of education driven by school-related sorting.

Table 6: Income Stratification Measures Based on Full Model With Neighborhood Sociodemographics

Panel A: Sample					Over-Exposure to Same Income Group	Percent Reduction in Over-Exposure
<i>Neighborhood Income Distribution</i>	<i>Household Income Quintile</i>					
	1-2	3	4-5	Overall		
1st/2nd Quintile	47.0	39.8	33.1	40.0	17.4%	
3rd Quintile	19.9	20.9	19.7	20.0		
4th/5th Quintile	33.1	39.3	47.3	40.0	18.1%	
Panel B: Pre-Simulation (Model's Prediction)					Over-Exposure to Same Income Group	Percent Reduction in Over-Exposure
<i>Neighborhood Income Distribution</i>	<i>Household Income Quintile</i>					
	1-2	3	4-5	Overall		
1st/2nd Quintile	46.0	41.0	33.5	40.0	15.0%	
3rd Quintile	20.5	20.5	19.3	20.0		
4th/5th Quintile	33.5	38.6	47.2	40.0	18.1%	
Panel C: Post-Simulation, Partial Equilibrium					Over-Exposure to Same Income Group	Percent Reduction in Over-Exposure
<i>Neighborhood Income Distribution</i>	<i>Household Income Quintile</i>					
	1-2	3	4-5	Overall		
1st/2nd Quintile	45.4	40.8	34.2	40.0	13.4%	10.5%
3rd Quintile	20.4	20.4	19.4	20.0		
4th/5th Quintile	34.2	38.9	46.3	40.0	15.8%	12.4%
Panel D: Post-Simulation, General Equilibrium					Over-Exposure to Same Income Group	Percent Reduction in Over-Exposure
<i>Neighborhood Income Distribution</i>	<i>Household Income Quintile</i>					
	1-2	3	4-5	Overall		
1st/2nd Quintile	44.6	40.6	35.1	40.0	11.5%	23.7%
3rd Quintile	20.3	20.3	19.6	20.0		
4th/5th Quintile	35.1	39.1	45.3	40.0	13.3%	26.5%

Notes (Columns 1-4): The first four columns of each panel of the table report the average neighborhood income distribution for households with the characteristics shown in the column heading. The first column, for example, shows the average income distribution in the neighborhoods in which households below the 40th percentile of the income distribution reside. Panel A reports these measures for the sample; Panel B reports predicted measures using the estimated parameters of the model; Panels C and D report measures based on a counterfactual simulation that sets preferences for the average test score identically to zero for all households. Panel C gives partial equilibrium measures that account for only the initial re-sorting; Panel D reports general equilibrium measures, which are derived from a new equilibrium in which households are allowed to re-sort on the basis of the new neighborhood sociodemographic compositions that follow the initial re-sorting induced by the elimination of preferences for school quality.

Notes (Column 5-6): The fifth column of the table reports the extent to which households in the bottom and top 40 percent of the income distribution or more are 'over-exposed' to other households in the same part of the income distribution respectively. So, for example, the first entry of this column is formed as $(47.0-40.0)/40.0$. The final column shows the percentage reduction in these measures in the simulation relative to the model's predictions, thus providing measures of the portion of neighborhood stratification on the basis of income driven by school-related sorting.

Appendix Table 1: Choice-Specific Constant Regressions

Sample	Without Neighborhood Sociodemographics			With Neighborhood Sociodemographics		
	full sample	within .25 mile of boundaries		full sample	within .25 mile of boundaries	
Boundary Fixed Effects	No	No	Yes	No	No	Yes
Observations	242,100	27,958	27,958	242,100	27,958	27,958
monthly housing price (/1000)	-10.23 (1.39)	-9.73 (1.13)	-11.34 (1.36)	-15.94 (1.71)	-15.97 (1.56)	-16.19 (1.69)
average test score (in standard deviations)	1.29 (0.02)	1.20 (0.05)	0.92 (0.01)	0.32 (0.03)	0.32 (0.08)	0.42 (0.01)
1 if unit owned	2.15 (0.03)	1.74 (0.09)	2.09 (0.01)	2.64 (0.05)	2.41 (0.14)	2.61 (0.01)
number of rooms	1.52 (0.02)	1.45 (0.04)	1.57 (0.01)	1.95 (0.02)	1.93 (0.07)	1.93 (0.01)
1 if built in 1980s	1.33 (0.04)	0.73 (0.11)	1.20 (0.02)	1.59 (0.06)	1.37 (0.17)	1.55 (0.02)
1 if built in 1960s or 1970s	0.29 (0.03)	0.09 (0.08)	0.17 (0.01)	0.22 (0.04)	0.12 (0.12)	0.07 (0.01)
elevation (/100)	0.22 (0.01)	-0.05 (0.02)	0.53 (0.01)	-0.02 (0.01)	-0.29 (0.04)	0.21 (0.01)
population density	-1.03 (0.04)	-1.49 (0.15)	-1.51 (0.03)	0.31 (0.07)	0.67 (0.25)	0.49 (0.03)
crime index	-0.03 (0.00)	-0.02 (0.01)	0.02 (0.00)	0.00 (0.00)	-0.02 (0.01)	0.03 (0.00)
% Census block group black				-5.18 (0.16)	-5.09 (0.51)	-4.32 (0.06)
% Census block group Hispanic				-0.07 (0.23)	0.29 (0.75)	2.25 (0.10)
% Census block group Asian				-1.55 (0.18)	-1.54 (0.60)	2.51 (0.09)
% block group college degree or more				4.56 (0.17)	3.29 (0.49)	2.23 (0.07)
average block group income				1.39 (0.02)	1.53 (0.06)	1.42 (0.01)
F-statistic for boundary fixed effects			4.545			3.963

Note: Specifications shown in the table also include controls for land use (% industrial, % residential, % commercial, % open space, % other) in 1, 2, 3, 4, and 5 mile rings around location and six variables that characterize the housing stock in each of these

Appendix Table 2. Interaction Parameter Estimates - Model Without Neighborhood Sociodemographics

	Average Test Score (+1 s.d.)	House Characteristics					Neighborhood Attributes			Distance to Work
		Monthly House Price (/1000)	Owner Occupied	Number of Rooms	Built in 1980s	Built in 1960-1979	Elevation (/100)	Population Density	Crime Index (0.001)	
Household Characteristics										
household income (/10,000)	0.050 (0.004)	0.121 (0.003)	0.305 (0.010)	0.074 (0.002)	0.142 (0.011)	0.038 (0.009)	0.016 (0.001)	0.028 (0.013)	-0.001 (0.001)	-0.004 (0.001)
1 if children under 18 in household	-0.190 (0.047)	0.063 (0.065)	-0.102 (0.094)	0.544 (0.025)	-0.316 (0.112)	0.146 (0.083)	0.010 (0.022)	-0.740 (0.101)	0.015 (0.005)	0.036 (0.005)
1 if black	-1.395 (0.080)	-0.941 (0.127)	-0.510 (0.167)	0.152 (0.044)	0.004 (0.211)	0.401 (0.144)	-0.062 (0.041)	-1.285 (0.159)	0.110 (0.007)	-0.023 (0.011)
1 if Hispanic	-0.642 (0.072)	0.168 (0.122)	-0.036 (0.130)	-0.268 (0.036)	-0.180 (0.164)	-0.157 (0.115)	-0.104 (0.040)	-0.155 (0.136)	0.050 (0.007)	0.014 (0.007)
1 if Asian	-0.167 (0.062)	0.315 (0.080)	1.765 (0.122)	-0.503 (0.031)	1.037 (0.145)	0.686 (0.108)	-0.015 (0.028)	0.941 (0.095)	0.030 (0.006)	0.003 (0.007)
1 if college degree or more	0.787 (0.053)	0.917 (0.071)	-0.032 (0.108)	-0.012 (0.029)	0.489 (0.135)	-0.045 (0.093)	0.225 (0.024)	-0.007 (0.111)	0.031 (0.006)	-0.006 (0.006)
1 if working	0.007 (0.049)	0.244 (0.067)	0.563 (0.103)	0.032 (0.027)	0.641 (0.125)	0.406 (0.086)	-0.048 (0.025)	-0.437 (0.097)	-0.027 (0.005)	-0.858 (0.008)
age (years)	0.015 (0.001)	0.010 (0.002)	0.090 (0.003)	0.004 (0.001)	-0.034 (0.004)	-0.009 (0.003)	0.003 (0.001)	-0.006 (0.003)	0.001 (0.000)	-0.001 (0.000)

Note: The parameters shown describe the elements of the utility function that interact household characteristics, shown in row headings, with choice characteristics, shown in column headings. Standard errors are in parentheses.

Appendix Table 3. Interaction Parameter Estimates - Model With Neighborhood Sociodemographics

	Average Test Score (+1 s.d.)	House Characteristics					Neighborhood Attributes			Neighborhood Sociodemographics					Distance to Work (miles)
		Monthly House Price (/1000)	Owner Occupied	Number of Rooms	Built in 1980s	Built in 1960-1979	Elevation (/100)	Population Density	Crime Index	% Block Group Black	% Block Group Hisp	% Block Group Asian	% Blk Group College	Blk Group Avg Income	
Household Characteristics															
household income (+10,000)	0.020 (0.005)	0.121 (0.004)	0.303 (0.011)	0.076 (0.003)	0.144 (0.012)	0.028 (0.009)	0.010 (0.002)	0.011 (0.017)	-0.001 (0.001)	-0.223 (0.060)	0.113 (0.064)	-0.009 (0.039)	0.385 (0.034)	0.012 (0.002)	-0.004 (0.001)
1 if children under 18 in household	0.102 (0.058)	0.231 (0.075)	-0.238 (0.103)	0.582 (0.028)	-0.399 (0.122)	0.095 (0.092)	0.051 (0.025)	-0.947 (0.127)	0.002 (0.006)	1.594 (0.416)	2.294 (0.527)	1.857 (0.387)	-2.171 (0.331)	0.055 (0.016)	0.027 (0.005)
1 if black	-0.282 (0.116)	0.143 (0.170)	-1.006 (0.205)	0.002 (0.053)	0.027 (0.253)	0.577 (0.184)	-0.068 (0.052)	-1.106 (0.228)	0.045 (0.009)	14.874 (0.560)	7.082 (0.888)	7.371 (0.747)	2.607 (0.680)	-0.023 (0.035)	-0.010 (0.013)
1 if Hispanic	-0.077 (0.089)	0.204 (0.139)	-0.138 (0.147)	-0.246 (0.041)	-0.147 (0.185)	-0.248 (0.131)	-0.067 (0.045)	-0.128 (0.169)	0.005 (0.008)	4.435 (0.568)	12.471 (0.620)	2.757 (0.587)	0.830 (0.492)	0.011 (0.022)	0.012 (0.008)
1 if Asian	0.072 (0.078)	0.558 (0.095)	1.633 (0.138)	-0.571 (0.035)	0.612 (0.166)	0.457 (0.123)	-0.006 (0.033)	-0.053 (0.132)	0.006 (0.007)	4.236 (0.562)	3.330 (0.721)	14.060 (0.429)	-0.016 (0.449)	-0.022 (0.022)	0.012 (0.007)
1 if college degree or more	0.200 (0.065)	0.501 (0.079)	0.428 (0.118)	0.006 (0.032)	0.588 (0.148)	0.106 (0.101)	0.031 (0.027)	0.486 (0.134)	0.022 (0.007)	1.279 (0.504)	-0.638 (0.607)	-1.935 (0.450)	8.986 (0.366)	0.009 (0.020)	0.009 (0.007)
1 if working	0.093 (0.062)	0.272 (0.074)	0.604 (0.113)	0.021 (0.030)	0.897 (0.138)	0.425 (0.096)	0.023 (0.028)	-0.515 (0.125)	-0.019 (0.007)	-0.712 (0.444)	-0.335 (0.563)	-0.434 (0.436)	-1.931 (0.350)	0.033 (0.016)	-0.896 (0.009)
age (years)	0.013 (0.002)	0.011 (0.002)	0.097 (0.003)	0.003 (0.001)	-0.033 (0.004)	-0.010 (0.003)	0.003 (0.001)	-0.011 (0.003)	0.001 (0.000)	-0.022 (0.013)	-0.085 (0.016)	-0.005 (0.013)	-0.018 (0.010)	0.001 (0.001)	-0.001 (0.000)

Note: The parameters shown describe the elements of the utility function that interact household characteristics, shown in row headings, with choice characteristics, shown in column headings. Standard errors are in parentheses.

APPENDICES FOR “Tiebout Sorting, Social Multipliers and the Demand for School Quality” by Patrick Bayer, Fernando Ferreira, and Robert McMillan

This document contains two appendices for the paper “Tiebout Sorting, Social Multipliers and the Demand for School Quality.” A *Data Appendix* documents the sources for the data and the construction of variables used in the analysis. An *Estimation Appendix* provides a full description of the estimation procedure used in our analysis, discussing a number of important details that we abstract from in the main text in to streamline that presentation.

DATA APPENDIX

1. Census Variables

House Prices. Because house values are self-reported, it is difficult to ascertain whether these prices represent the current market value of the property, especially if the owner purchased the house many years earlier. Fortunately, the Census contains other information that helps us to examine this issue and correct house values accordingly. In particular, the Census asks owners to report a continuous measure of their annual property tax payment. The rules associated with Proposition 13 imply that the vast majority of property tax payments in California should represent exactly 1 percent of the transaction price of the house at the time the current owner bought the property or the value of the house in 1978. Thus, by combining information about property tax payments and the year that the owner bought the house (also provided in the Census in relatively small ranges), we are able to construct a measure of the rate of appreciation implied by each household’s self-reported house value. We use this information to modify house values for those individuals who report values much closer to the original transaction price rather than current market value. In our study most households list the purchase price of their house rather than an estimated market value for their house. Thus if two identical houses were found in the census data but one was last sold in 1989 and one was last sold in 1969 we find on average the listed market price of the more recently sold house is on average 15 percent higher than the other house.

A second deficiency of the house values reported in the Census is that they are top-coded at \$500,000, a top-code that is often binding in California. Again, because the property tax payment variable is continuous and not top-coded, it provides information useful in distinguishing the values of the upper tail of the value distribution. We find that top-coding was fairly predominant in the Bay Area and that higher top-codes may be useful to gain a better understanding of house prices in expensive markets like California or New York.

The exact procedure that we use to adjust self-reported house values is as follows. We first regress the log of self-reported house value on the log of the estimated transaction price (100 times the property tax payment), and a series of dummy variables that characterize the tenure of the current owner:

$$(A1) \quad \log(V_j) = \alpha_1 \log(T_j) + \alpha_2 y_j + w_j$$

where V_j represents the self-reported house value, T_j represents the estimated transaction price, and y_j represents a series of dummy variables for the year that the owner bought the house. If owner-estimated house values were indeed current market values and houses were identical except for owner tenure, this regression would return an estimate of 1 for α_1 and the estimated α_2 coefficients would indicate the appreciation of house values in the Bay Area over the full period of analysis. If owners tend to underreport house values, especially when they have lived in the house for a long time, the estimated α_2 parameters will likewise underreport appreciation in the market. In this way, the estimated α_2 parameters represent a conservative estimate of appreciation. Given the estimates of equation (2), we construct a predicted house value for each house in the sample and replace the owner-reported value with this measure when this predicted measure exceeds the owner-reported value. In practice, in order to allow for different rates of appreciation in different regions of the housing market, we conduct these regressions separately for each of the 45 Census PUMA (areas with at least 100,000 people) in our sample and allow appreciation to vary with a small set of house characteristics within each PUMA. In this way, the first adjustment that we make to house prices is to adjust owner-reported values for likely under-reporting.

The adjustment to top-coded house prices uses the same approach, using the information on property taxes that are continuous and not top-coded. Using estimates of equation (2) based on a sample of

houses that does not include the top-coded house values, we construct predicted house values for all top-coded houses. This allows us to assign continuous house values for top-coded measures.

Reported Rental Value. We next examined questions of reported monthly rents. While rents are presumably not subject to the same degree of misreporting as house values, it is still the case that renters who have occupied a unit for a long period of time generally receive some form of tenure discount. In some cases, this tenure discount may arise from explicit rent control, but implicit tenure discounts generally occur in rental markets even when the property is not subject to formal rent control. Thus while, this will not lead to errors in the answering of the listed census question it may lead to an inaccurate comparison of rents faced by households if they needed to move. In order to get a more accurate measure of the market rent for each rental unit, we utilize a series of locally based hedonic price regressions in order to estimate the discount associated with different durations of tenure in each of over 40 sub-regions within the Bay Area.

In order to get a better estimate of market rents for each renter-occupied unit in our sample, we regress the log of reported rent R_j on a series of dummy variables that characterize the tenure of the current renter, y_j , as well as a series of variables that characterize other features of the house and neighborhood X_j :

$$(A2) \quad \log(R_j) = \mathbf{b}_1 y_j + \mathbf{b}_2 X_j + \mathbf{u}_j$$

again running these regressions separately for each of the 45 PUMAs in our sample. To the extent that the additional house and neighborhood variables included in equation (3) control for differences between the stock of rental units with long-term vs. short-term tenants, the β_1 parameters provide an estimate of the tenure discount in each PUMA.¹ In order to construct estimates of market rents for each rental unit in our sample, then, we inflate rents based on the length of time that the household has occupied the unit using the estimates of β_1 from equation (2). In this way, these three price adjustments bring the measures for rents and house values reported in the Census reasonably close to market rates.

Calculating Cost Per Unit of Housing Across Tenure Status. Finally, in order to make owner- and renter-occupied housing prices comparable in our analysis we need to calculate a current rental value for housing. Because house prices reflect the expectations about the future rents for the property they incorporate beliefs about future housing appreciation. To appropriately deflate housing values – and especially to control for differences in expectations about appreciation in different segments of the Bay Area housing market – we regress the log of house price (whether monthly rent or house value) Π_j on an indicator for whether the housing unit is owner-occupied o_j and a series of additional controls for features of the house including the number of rooms, number of bedrooms, types of structure (single-family detached, unit in various sized buildings, etc.), and age of the housing structure as well as a series of neighborhood controls X_j :

$$(A3) \quad \log(\Pi_j) = \mathbf{g}_1 o_j + \mathbf{g}_2 X_j + \mathbf{h}_j$$

We estimate these hedonic price regressions for each of 40 sub-regions (Census Public Use Microdata Areas - PUMAs) of the Bay Area housing market. These regressions return an estimate of the ratio of house values to rents for each of these sub-regions and we use these ratios to convert house values to a measure of current monthly rent.

2. External Data

We next discuss the additional variables we have added to the Census data to provide a more nuanced understanding of the neighborhood characteristics that affect house prices and residential location decisions. These data sets are linked to census blocks and can be used to determine the appropriateness of the questions and sampling techniques used. This additional data includes:

¹ Interestingly, while we estimate tenure discounts in all PUMAs, the estimated tenure discounts are substantially greater for rental units in San Francisco and Berkeley, the two largest jurisdictions in the Bay Area that had formal rent control in 1990.

School and School District Data. The Teale data center in California provided a crosswalk that matches all Census blocks in California to the corresponding public school district. We have further matched Census blocks to particular schools using a variety of procedures that takes account of the location (at the block level) of each Census block within a school district and the precise location of schools within the district using information on location from the Department of Education. Other school information in these data include:

- 1992-93 CLAS dataset provides detailed information about school performance and peer group measures. The CLAS was a test administered in the early 1990s that will give us information on student performance in math, literature and writing for grades 4, 8 and 10. This dataset presents information on student characteristics and grades for students at each school overall and across different classifications of students, including by race and education of parents.
- 1991-2 CBEDS (California Board of Education data sets) datasets including information from the SIF (school information form) which includes information on the ethnic/racial and gender make-up of students, PAIF – which is a teacher based form that provides detailed information about teacher experience, education and certification backgrounds and information on the classes each teacher teaches, and (LEP census) a language census that provides information on the languages spoken by limited-English speaking students.

Procedures for Assigning School Data. While we have an exact assignment of Census blocks to school districts, we have only been able to attain precise maps that describe the way that city blocks are assigned to schools in 1990 for Alameda County. In the absence of information about within-district school attendance areas, we employ the alternative approaches for linking each house to a school. The crudest procedure assigns average school district characteristics to every house falling in the school district. A refinement on this makes use of distance-weighted averages. For a house in a given Census block, we calculate the distance between that Census block and each school in the school district. We have detailed information characterizing each school and construct weighted averages of each school characteristic, weighting by the reciprocal of the distance-squared as well as enrollment.

As a third approach we simply assign each house to the closest school within the appropriate school district. Our preferred approach (which we use for the results reported in the paper) refines this closest-school assignment by using information about individual children living in each Census block - their age and whether they are enrolled in public school. In particular, we modify the closest-school assignment technique by attempting to match the observed fourth grade enrollment for every school in every school district in the Bay Area. Adjusting for the sampling implicit in the long form of the Census, the 'true' assignment of houses to schools must give rise to the overall fourth grade enrollments observed in the data.

These aggregate numbers provide the basis for the following intuitive procedure: we begin by calculating the five closest schools to each Census block. As an initial assignment, each Census block and all the fourth graders in it are assigned to the closest school. We then calculate the total predicted enrollment in each school, and compare this with the actual enrollment. If a school has excess demand, we reassign Census blocks out of its catchment area, while if a school has excess supply, we expand the school's catchment area to include more districts.

To carry out this adjustment, we rank schools on the basis of the (absolute value of) their prediction error, dealing with the schools that have the greatest excess demand/supply first. If the school has excess demand, we reassign the Census block that has the closest second school (recalling that we record the five closest schools to each Census block, in order), as long as that second school has excess supply. If a school has excess supply, we reassign to it the closest school district currently assigned to a school with excess demand. We make gradual adjustments, reassigning one Census block from each school in disequilibrium each iteration. This gradual adjustment of assignments of Census blocks to schools continues until we have 'market clearing' (within a certain tolerance) for each school. Our actual algorithm converges quickly and produces plausible adjustments to the initial, closest-school assignment.

Land use. Information on land use/land cover digital data is collected by USGS and converted to ARC/INFO by the EPA available at: <http://www.epa.gov/ost/basins/> for 1988. We have calculated for each Census block, the percentage of land in a 1/4, 1/2, 1, 2, 3, 4 and 5-mile radii that is used for commercial,

residential, industrial, forest (including parks), water (lakes, beaches, reservoirs), urban (mixed urban or built up), transportation (roads, railroad tracks, utilities) and other uses.

Crime data. Information on crime was drawn from the rankings of zipcodes on a scale of 1-10 on the risk of violent crime (homicide, rape or robbery). A score of 5 is the average risk of violent crime and a score of 1 indicates a risk 1/5 the national average and a 10 is 10 or more times the national average. These ratings are provided by CAP index and were downloaded from APBNews.com.

Geography and Topography. The Teale data center in California provided information on the elevation, latitude and longitude of each Census block.

ESTIMATION APPENDIX

Estimation of the model follows a two-stage procedure closely related to that developed in Berry, Levinsohn, and Pakes (1995). It is helpful in describing the estimation procedure to first introduce some notation. In particular, we rewrite the indirect utility function as:

$$(A4) \quad V_h^i = \mathbf{d}_h + \mathbf{I}_h^i + \mathbf{e}_h^i$$

where

$$(A5) \quad \mathbf{d}_h = \mathbf{a}_{0X} X_h + \mathbf{a}_{0Z} Z_h - \mathbf{a}_{0p} p_h + \mathbf{x}_h$$

and

$$(A6) \quad \mathbf{I}_h^i = \left(\sum_{k=1}^K \mathbf{a}_{kX} z_k^i \right) X_h + \left(\sum_{k=1}^K \mathbf{a}_{kZ} z_k^i \right) Z_h - \left(\sum_{k=1}^K \mathbf{a}_{kp} z_k^i \right) p_h - \left(\sum_{k=1}^K \mathbf{a}_{kd} z_k^i \right) \mathbf{d}_h.$$

In equation (A5), \mathbf{d}_h captures the portion of utility provided by housing type h that is common to all households, and in (A6), k indexes household characteristics. When the household characteristics included in the model are constructed to have mean zero, \mathbf{d}_h is the *mean indirect utility* provided by housing choice h . The unobservable component of \mathbf{d}_h , \mathbf{x}_h , captures the portion of unobserved preferences for housing choice h that is correlated across households, while \mathbf{e}_h^i represents unobserved preferences over and above this shared component.

The estimator is a two-stage procedure. The first stage selects the heterogeneous parameters \mathbf{I}_h and mean indirect utilities \mathbf{d}_h that maximize the probability that the model correctly predicts each individual's location decision conditional on the full set of observed housing and neighborhood attributes, including those endogenously determined. Formally, the validity of this first stage requires two assumptions: that the observed location decisions are individually optimal, given the collective choices made by other households and the vector of market-clearing prices, *and* that households are sufficiently small such that they do not interact strategically with respect to particular draws on \mathbf{e} . This latter assumption ensures that households can each effectively integrate out the idiosyncratic preferences of all others when making their own location decisions and so that no household's particular idiosyncratic preferences affect the equilibrium. Thus the vector of idiosyncratic preferences \mathbf{e} is uncorrelated with the prices and neighborhood sociodemographic characteristics that arise in *any* equilibrium.

In essence, the first-stage of the estimation procedure is equivalent to a Maximum Likelihood procedure that treats housing prices and neighborhood sociodemographic characteristics as exogenous from the individual's point-of-view. Importantly, the assumption that prices and neighborhood sociodemographic characteristics are uncorrelated with the vector of idiosyncratic preferences \mathbf{e} does not imply that they are uncorrelated with the full error term, as we explicitly allow for a portion of unobserved preferences, \mathbf{x} that is correlated with price and endogenous neighborhood characteristics in equilibrium. This correlation is addressed in the second stage of the estimation procedure, in which the vector \mathbf{d} estimated in the first stage is decomposed into components.

Operationally, for any combination of the heterogeneous parameters in \mathbf{I} and mean indirect utilities, \mathbf{d} , the model predicts the probability that each household i chooses house type h . We assume that \mathbf{e}_h^i is drawn from the extreme value distribution, in which case this probability can be written:

$$(A7) \quad P_h^i = \frac{\exp(\mathbf{d}_h + \hat{I}_h^i)}{\sum_k \exp(\mathbf{d}_k + \hat{I}_k^i)}$$

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

$$(A8) \quad \tilde{\ell} = \sum_i \sum_h I_h^i \ln(P_h^i)$$

where I_h^i 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 \mathbf{I} and the vector of mean indirect utilities to maximize $\tilde{\ell}$. Notice that the quasi-likelihood function developed here is based solely on the notion that each household's residential location is optimal given the set of observed prices and the location decisions of other households.

The Mechanics of the First Stage of the Estimation. Intuitively, it is easy to see how this first stage of the estimation procedure ties down the heterogeneous parameters – those involving an interaction of household characteristics with housing and neighborhood characteristics. If more educated households are more likely to choose houses near better schools in the data for instance, a positive interaction of education and school quality will allow the model to fit the data better than a negative interaction would. What is less intuitive is the way the vector of mean indirect utilities is determined. To better understand the mechanics of the first stage of the estimation, it is helpful to write the first-order conditions related to \mathbf{d} :

$$(A9) \quad \frac{\partial \tilde{\ell}}{\partial \mathbf{d}_h} = \sum_{i \in h} \frac{\partial \ln(P_h^i)}{\partial \mathbf{d}_h} + \sum_{i \notin h} \frac{\partial \ln(P_h^i)}{\partial \mathbf{d}_h} = \sum_{i \in h} (1 - P_h^i) + \sum_{i \notin h} (-P_h^i) = S_h - \sum_i (P_h^i) = 0$$

It is apparent that the quasi-likelihood function is maximized at the vector \mathbf{d} that forces the sum of the probabilities that each observed individual chooses each house type to equal the total supply of such houses: $\sum_i (P_h^i) = S_h \forall h$. That this condition must hold for all house types results from a fundamental

trade-off in $\tilde{\ell}$. In particular, an increase in any \mathbf{d} raises the probability that each household in the sample chooses house type h . While this increases the probability that the model correctly predicts the choice of the households that actually reside in houses of type h , it decreases the probability that all of the other households in the sample make the correct choice. Thus the first stage of the estimation procedure consists of choosing the interaction parameters that best match each individual with their chosen house, while ensuring that total predicted demand equals supply for each house type.

For any set of interaction parameters (those in λ), a contraction mapping can be used to calculate the vector \mathbf{d} that solves the set of first order conditions: $\sum_i (P_h^i) = S_h \forall h$. For our application, the contraction mapping is simply:

$$(A10) \quad \mathbf{d}_h^{t+1} = \mathbf{d}_h^t - \ln \left(\sum_i \hat{P}_h^i / S_h \right)$$

where t indexes the iterations of the contraction mapping. Using this contraction mapping, it is possible to solve quickly for an estimate of the full vector $\hat{\mathbf{d}}$ even when it contains a large number of elements, thereby dramatically reducing the computational burden in the first stage of the estimation procedure.²

Notice that while we have not explicitly enforced the market clearing conditions derived above, the conditions that result from maximizing the quasi-likelihood function with respect to \mathbf{d} are identical to the market-clearing conditions shown in equation (6) in the main text. Thus, there is a clear duality between the equilibrating role of prices in our characterization of equilibrium in the housing market and the way that the vector of mean indirect utilities is determined as a result of maximizing the likelihood that each household chooses its appropriate house conditional on prices and housing and neighborhood attributes.

The Second Stage. Having estimated the vector of mean indirect utilities in the first stage of the estimation procedure, the second stage involves decomposing \mathbf{d} into observable and unobservable components according to equation (A5).³ Because households sort across locations based in part on the portion of housing and neighborhood quality unobserved by the researcher, housing prices, average test scores, and neighborhood sociodemographic characteristics are almost certainly correlated with \mathbf{x} and consequently the corresponding endogeneity problems must be confronted. We describe the use of boundary fixed effects to address the endogeneity of average test scores and neighborhood sociodemographic characteristics in the main text.

To deal with the correlation of price and unobserved housing/neighborhood quality, \mathbf{x} , we instrument for price. The particular instrument that we develop takes advantage of an inherent feature of housing markets: that the demand for a house in a particular neighborhood is affected not only by the features of the neighborhood itself but also by the availability of alternative houses and neighborhoods in the wider region. For example, neighborhoods that possess certain amenities that are unique or difficult to replicate will command higher prices in equilibrium, partly because of this scarcity. The exogenous attributes of houses and neighborhoods at a reasonable distance from a particular neighborhood serve as suitable instruments for price, as the attributes of these more distant neighborhoods affect equilibrium prices but not the utility derived from living in the neighborhood.⁴

In practice, the precision of the estimation is improved significantly when the logic of this IV strategy is used to construct a single variable that approximates the optimal instrument. In particular, we construct an instrument by solving for the vector of prices that would clear the market when only exogenous features of houses and neighborhoods are included in the utility function. This instrument captures the portion of housing price variation attributable to the distribution of the exogenous features of houses and neighborhoods throughout the region, summarizing this information in a single variable.

A couple of additional practical items are worth describing. First, the construction of the instrument requires an initial conjecture as to the parameters associated with exogenous housing and neighborhood attributes. We obtain such an initial conjecture for the parameters of the mean indirect utility equation by making a reasonable guess as to the price coefficient and then estimating equation (11) via OLS, bringing the price term to the left hand side of the equation. Using the resulting coefficients on \mathbf{X} from this regression along with those obtained in the first stage, we then calculate the vector of housing prices that clears the market, $\hat{\mathbf{p}}^*(\mathbf{X}_h, \mathbf{Z}^i)$, setting $\mathbf{x}_h=0$ for all h , and including only *exogenous* choice

² It is worth emphasizing that a separate vector \mathbf{d} is calculated for each set of interaction parameters – and at the optimum, this procedure returns the quasi-ML estimates of the interaction parameters and the vector of mean indirect utilities \mathbf{d} .

³ Notice that the set of observed residential choices provides no information that distinguishes the components of \mathbf{d} . That is, regardless of the way \mathbf{d} is broken into components, the effect on choice probabilities is the same.

⁴ Put another way, for most individuals, the relevant extent of the housing market is much larger when they are searching for a house (they might live, for example, to the north, south, east, or west of their job location) than when they actually choose a residence, in which case the characteristics of houses on the opposite side of town likely have only a minimal direct impact on utility. It is the fact that a much broader set of houses is in play during the search process that implies that the characteristics of the housing stock on the other side of town will influence equilibrium prices.

characteristics in the model.⁵ In the results reported below, we include a full set of controls for the characteristics of a house and its neighborhood as well as five variables that describe land use⁶ and six variables that describe the housing stock⁷ in each of the 1, 2, 3, 4, and 5 mile rings around the house. In this way, the additional information embedded in our instrument derives from the exogenous features of the housing stock and land use in a region beyond five miles from the house in question.⁸

Characterizing the Housing Market – A Practical Issue. A final practical issue for estimation concerns the way the choices that characterize the housing market should be defined. This modeling decision essentially corresponds to an assumption regarding the way demand for particular houses in the market is determined. The trade-offs implicit in the required assumption can be seen using a simple example: Consider a city neighborhood with two types of housing structures, one of which is more prevalent than the other, with all houses in the neighborhood selling for the same price. To simplify this discussion, further assume that households have identical tastes. In this case, if we characterized the choice set as the two types of structure, we would infer that the more prevalent structure provided higher mean direct utility; this is necessary to explain why more households choose that structure given equal prices. If, on the other hand, we characterized the housing market by randomly drawing a subset of the houses in the neighborhood, we would infer that all of the houses in the neighborhood offered the same utility. We do not see any strong *a priori* for making one of these choices versus the other. Moreover, given that any definition of ‘type’ would be based only on the limited characteristics observed in the data, we adopt the second option described above, simply characterizing housing types as the 1-in-7 random sample of the houses observed in our Census dataset. This characterization also facilitates comparisons with the hedonic price regression literature; with this characterization of the choice set, a hedonic price regression corresponds to estimating mean preferences under the assumption of no heterogeneity in household tastes.⁹

Asymptotic Properties of the Estimator. As described in McFadden (1978), an attractive aspect of the underlying IIA property for each individual is that we can estimate the model using only a sample of the alternatives not selected by the individual. This permits estimation despite having many alternatives – i.e., many distinct house types. More generally, our problem fits within a class of models for which the asymptotic distribution theory has been developed. In this sub-section, we summarize the requirements necessary for the consistency and asymptotic normality of our estimates and provide some intuition for these conditions.

In general, there are three dimensions in which our sample can grow large: H (number of housing types), N (number of individuals in the sample), or C (number of non-chosen alternatives drawn for each individual). For any set of distinct housing alternatives of size H and any random sampling of these alternatives of size C , the consistency and asymptotic normality of the first-stage estimates (δ, \mathbf{q}_i) follows directly as long as N grows large. This is the central result of McFadden (1978), justifying the use of a random sample of the full census of alternatives. Intuitively, even if each household is assigned only one

⁵ To obtain the final estimates reported in the paper, we repeat this procedure using the estimated parameters from the initial estimation to construct a new price instrument for the next iteration. While using such an iterative process is not necessary to ensure consistency, in practice it ensures that the final estimates are not sensitive to our initial conjecture of the coefficient on price. For this reason, we believe that this iterative procedure is likely to be more efficient than applying the procedure once, but we do not have a proof of this.

⁶ That is: percent industrial, percent commercial, percent residential, percent open space, and percent other.

⁷ The housing stock variables are: percent owner-occupied single family homes with 7 rooms or more; percent owner-occupied single family homes with less than 7 rooms; percent renter-occupied single family homes; percent renter-occupied units in large apartment buildings; percent of units in small apartment buildings; percent other.

⁸ In first-stage price regressions, this instrument, which is derived entirely from the exogenous characteristics of the alternatives and the distribution of household characteristics in the population, adds significantly to the predictive power of these regressions. In each specification, the optimal price instrument is strongly predictive of price, over and above the set of variables included directly in \mathbf{X} , increasing the R^2 of each regression by approximately 4 percentage points.

⁹ Nothing theoretically prevents estimation of the model under an alternative assumption concerning housing choices.

randomly drawn alternative in addition to its own choice, the number of times that each house type is sampled (the dimension in which the choice-specific constants are identified) grows as a fixed fraction of N .

If the true vector \mathbf{d} were used in the second stage of the estimation procedure, the consistency and asymptotic normality of the second-stage estimates \mathbf{q}_i would follow as long as $H \rightarrow \infty$ ¹⁰. In practice, ensuring the consistency and asymptotic normality of the second-stage estimates is complicated by the fact the vector \mathbf{d} is estimated rather than known. Berry, Linton, and Pakes (2002) develop the asymptotic distribution theory for the second stage estimates \mathbf{q}_i for a broad class of models that contains our model as a special case, and consequently we employ their results. In particular, the consistency of the second-stage estimates follows as long as $H \rightarrow \infty$ and N grows fast enough relative to H such that $H \log H / N$ goes to zero, while asymptotic normality at rate \sqrt{H} follows as long as H^2 / N is bounded. Intuitively, these conditions ensure that the noise in the estimate of \mathbf{d} becomes inconsequential asymptotically and thus that the asymptotic distribution of \mathbf{q}_i is dominated by the randomness in \mathbf{x} as it would be if \mathbf{d} were known.

Given that the consistency and asymptotic normality of the second stage estimates requires the number of individuals in the sample to go to infinity at a faster rate than the number of distinct housing units, it is important to be clear about the implications of the way that we characterize the housing market in the paper. In particular, we characterize the set of available housing types using the 1-in-7 random sample of the housing units in the metropolitan area observed in our Census dataset. Superficially, this characterization seems to imply that the number of housing types is as great as the number of households in the sample, which appears at odds with the requirements for the establishing the key asymptotic properties of our model.

It is important to note, however, the housing market may be characterized by a much smaller sample of houses, with each ‘true’ house type showing up many times in our large sample. Consider, for example, using a large choice set of 250,000 housing units, when the market could be fully characterized by 25,000 ‘true’ house types, with each ‘true’ house type showing up an average of 10 times in the larger choice set. On the one hand, the 250,000 observations could be used to calculate the market share of each of the 25,000 ‘true’ house types, with market shares averaging $1/25,000$ and the second stage \mathbf{d} regressions based on 25,000 observations. On the other hand, separate market shares equal to $1/250,000$ could be attributed to each house observed in the larger sample and the second stage regression based on the larger sample of 250,000. These regressions would return exactly the same estimates, as the former regression is a direct aggregation of the latter. What is important from the point-of-view of the asymptotic properties of the model is not that the number of individuals increases faster than then number of housing choices used in the analysis, but rather that the number of individuals increases fast enough relative to the number of truly distinct housing types in the market. That the number of distinct housing types in the market grows at a rate slower than the number of households seems plausible.

¹⁰ This condition requires certain regularity conditions. See Berry, Linton, and Pakes (2002) for details.