CHOICE AND COMPETITION IN LOCAL EDUCATION MARKETS[*]

Patrick Bayer
Department of Economics
Duke University and NBER

Robert McMillan
Department of Economics
University of Toronto and NBER

April 20, 2009

Abstract

Prompted by widespread concerns about public school quality, a growing empirical literature has measured the effects of greater choice on school performance. This paper contributes to that literature in three ways, first by noting that the *overall* effect of greater choice – the focus of prior research – can be decomposed into demand and supply responses: knowing the relative sizes of the two is very relevant for policy. Second, using rich data from a large metropolitan area, the analysis generates a direct and intuitive measure of the competition each school faces in the form of a school-specific slope, capturing the extent to which reductions in school quality lead to reductions in demand. Third, the paper provides evidence that these competition measures are strongly related to school performance: a one standard deviation increase in the competitiveness of a school's local environment leads to a 0.1 standard deviation increase in average test scores, controlling for a host of other factors. This positive relationship is robust, and is consistent with quite strong supply responsiveness on the part of public schools, of relevance to the broader school choice debate.

Key words: School Competition, School Choice, Education Demand, School Performance.

JEL: I20, H41, R21

---

[*] This is a revised version of NBER Working Paper #11802. Thanks to Pat Bajari, Steve Berry, Tim Besley, Damon Clark, Dennis Epple, Fernando Ferreira, Caroline Hoxby, Brian Jacob, Steve Rivkin, Jesse Rothstein, Holger Sieg, Douglas Staiger, Chris Taber, Chris Timmins, and especially Tim Bresnahan and Tom Nechyba, for helpful discussions about this research. We would also like to thank seminar participants at Bristol University, the University of Houston, LSE, the NBER, Queen's, Rice, Texas A&M, and PPIC for additional comments. Melissa Knox provided excellent research assistance. We gratefully acknowledge financial support from the U.S. Department of Education, the National Science Foundation (grant SES-0137289), the Public Policy Institute of California, and SSHRC. The research in this paper was conducted while the authors were Special Sworn Status researchers of the U.S. Census Bureau, at the Berkeley and Triangle Census Research Data Centers. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau. This paper has been screened to ensure that no confidential data are revealed.

1 INTRODUCTION

Dissatisfaction with the quality of public education in North America and beyond has prompted considerable interest in reforms that increase choice and stimulate competition in the school system. The standard notion, drawing on the theory of the firm, is that increasing the range of options households can choose from forces incumbent public schools to compete more vigorously – for market share, better students, and financial resources. In turn, this should lead schools to better respond to the needs of students and their families, use available resources more efficiently, and exert greater effort, thereby improving public school performance. Because schools lack a strict profit motive, however, the extent to which public schools do actually respond to increased competition has become a keenly debated empirical question in policy and academic circles.

In prior research, numerous papers have examined the effects of greater choice on public school performance, typically using across-metropolitan area variation in indices of public school concentration or measures of private school availability.[1] This literature has grappled with challenging endogeneity problems related to the simultaneous determination of public school performance and such measures of choice, in addition to omitted variable biases.[2] In part because of non-trivial measurement difficulties, the evidence to emerge from the literature is mixed, ranging from little effect to a modest positive impact of increased choice on school performance.[3]

While the concepts of choice and competition are invariably linked in the broad policy and academic debates about education reform, they are not synonymous from an economic perspective.[4] In this setting, choice relates to the availability of schooling alternatives open to households, while the economic notion of competition is best captured by the slope of the quality demand curve that a given school faces. This slope measures the extent to which a school's enrollment and resources are affected by a change in its performance – the steeper the slope, the more competitive is the school's environment. In considering the effect of *competition* on school

---

[1] See Belfield and Levin (2002) for an extensive review of forty one empirical studies in this literature.

[2] See Hoxby (1994, 2000), Figlio and Stone (1999), Jepsen (1999), Hanushek and Rivkin (2003), and Rothstein (2006).

[3] Researchers have also examined the efficiency of public school spending using alternative research designs. Barrow and Rouse (2004), for example, examine the efficiency of per pupil spending using variation in state aid while Millimet and Rangaprasad (2005) explore the spatial autocorrelation in school input decisions.

[4] In principle, choice could increase without any increase in competition, and vice-versa. For example, private school vouchers could be given to parents wholeheartedly opposed to private schools; likewise, the competitiveness of the education market place could change without any change in the set of available schooling options (choice), for instance through improvements in the quality of information available to households about schools and their performance.

performance, one would ideally like to isolate the impact of a change in this primitive – the slope of a school's quality demand curve – on school performance.

In practice, the measures of choice used in the literature tend to be related to competition because an increase in the availability of schooling options or in the ease of selecting these options is likely to steepen the slope of the demand curve faced by incumbent schools, thereby making the market more competitive. We note, however, that the overall impact of increased choice on school performance depends on both demand and supply factors, combining (i) the effect of increased choice on the competitiveness of local market (i.e., extent to which increased choice steepens the slope of the demand curve faced by schools), and (ii) the effect of increased competition on school performance. The first of these, naturally, reflects factors on the demand side of the education market, and is likely to vary depending on the way that choice increases.[5] The latter effect isolates the supply-side response to greater competition – the response of individual schools to variation in the competitiveness of the environment they operate in.

By focusing on the overall effect of *choice* on school performance, the existing literature provides only indirect guidance as to the magnitude of the responsiveness of schools to *competition* (supply responsiveness). Yet the distinction between demand and supply responsiveness is important both for understanding the economics of education markets and for policy. To the extent that an overall modest positive effect of choice on performance is driven by a sizeable impact of choice on competition and a limited responsiveness of schools to increased competition, this suggests that policies targeted at improving incentives for schools to respond to competition are likely to have strong performance effects. In contrast, if schools are already responsive to changes in competition (as we will find) but the current system does not engender a very competitive environment, policies aimed at increasing demand responsiveness would seem more promising. Such policies might, for example, increase the ability of parents to choose from a wider set of schools or provide better information about quality differences among the available schooling options.

In terms of magnitudes, a reading of the existing education demand literature suggests that the demand-side component of the overall effect of choice on school performance may be quite small. This conclusion is robust across a wide variety of different research designs: (i) examining differences in housing prices along school assignment boundaries or as school assignments are changed over time;[6] (ii) estimating the heterogeneity in willingness-to-pay for

---

[5] For example, choice variation from geographic accessibility to private schools may have a quite different impact on competition compared to policies such as within-district school choice programs.

[6] 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

school quality using a model of residential sorting;[7] (iii) examining the academic performance of winners versus losers following randomized school choice lotteries;[8] and (iv) examining how student performance varies with measures of school district enrollment concentration across metropolitan areas.[9] This still leaves open the possibility that the responsiveness of schools to increases in competition may be sizeable, despite the mixed results of the previous literature exploring the impact of choice on performance.

In this paper, we set out evidence based on a new approach for studying the effects of competition on school performance. At the heart of the analysis, we develop a school-specific competition measure, which captures the extent to which reductions in school quality would lead to reductions in demand. This measure has intuitive appeal: a school is taken to operate in a competitive environment if slight reductions in school quality would lead to a significant reduction in demand in the neighborhoods from which the school draws its students.[10]

To generate these school-specific competition measures, we begin by estimating a rich model of residential sorting using data on the precise location decisions of nearly a quarter of a million households in the San Francisco Bay Area. In the course of estimating the model, we are careful to address important endogeneity problems that arise due to correlations of school quality and observable neighborhood sociodemographic characteristics with unobserved aspects of housing and neighborhood quality – correlations induced by residential sorting. Here, we implement an extended boundary fixed effects strategy building on Black's (1999) approach, along with instruments that account for the endogeneity of house prices. The resulting preference estimates incorporate a considerable amount of household-level heterogeneity, and are reasonable in magnitude across a wide set of housing, schooling, and neighborhood attributes.

We use these rich preference estimates in combination with an equilibrium model of sorting to estimate a measure of the slope of the demand curve faced by each of over 700 elementary schools in the Bay Area. In particular, we carry out a series of school-specific simulations that yield the predicted change in demand from raising that school's quality by a fixed amount while holding the quality of all other schools fixed. Our main analysis involves

---

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.

[7] Bayer, Ferreira, and McMillan (2007) 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.

[8] Cullen, Jacob, and Levitt (2006) 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.

[9] Rothstein (2006) finds little evidence that sorting is related to school rather than peer characteristics.

[10] One can measure demand in a variety of ways. We focus on changes in local house values.

exploring the relationship between these slope measures and public school performance, as measured by standardized tests. If schools simply maximize quality, as measured by test scores given their available resources, then the slope of the demand curve would be irrelevant to a school's quality-setting decision. In contrast, if schools were rent-seeking, then the slope would play a key role, as in the textbook theory of the firm.

Our results from this empirical exercise are striking and robust. We find clear evidence that greater competition (captured by a steeper slope of the quality demand curve) is associated with significant increases in public school achievement scores, with little effect on resource use, indicating that school productivity improves. These findings are robust, persisting when we control for a wide set of student, school, household and neighborhood attributes (including all the variables that are included in the demand model); and the same findings hold regardless of which student, school, and neighborhood characteristics we condition on. The magnitude of the estimated effect is quite large: a one standard deviation increase in the competitiveness of a school's local market within the Bay Area is associated with around a 0.1 standard deviation increase in school performance.

These achievement increases are accompanied by, if anything, slight reductions in teacher quality measures and other school inputs. This helps to allay concerns about the role of omitted variables in these regressions. Such concerns are also eased by specifications that demonstrate that the predicted effect of our slope measures on neighborhood capital income is negative and insignificant. In terms of heterogeneous effects, while we find that all types of schools respond to increased competition, school responsiveness is greater in more educated communities.

The impact of these measures prompts us to investigate the underlying sources of variation in the slopes. In general, our estimated slope measures will depend on two factors: (i) the set of households served by the school (how mobile the types of household are that a school serves), and (ii) the availability of close substitutes in the education market. By including controls for a wide set of student, school, household and neighborhood attributes, including all the variables that are included in the demand model in the analysis, we seek to isolate only the latter form of variation (i.e., the availability of close substitutes) in our competition measure. To provide additional confidence in our baseline results, we relate our competition measure to indices created directly from the data that characterize the availability of close substitutes.[11] In particular, we construct dissimilarity indices that describe the extent to which given schools and neighborhoods are isolated in quality space relative to their closest neighbors. We demonstrate that our competition measure is strongly negatively correlated with these dissimilarity measures: the more dissimilar a school is from its neighbors, so its demand is less responsive to quality changes. In turn, these dissimilarity measures are correlated with school performance in a manner consistent with our baseline results.

These findings have implications for the broader choice debate. That we find sizeable effects in the San Francisco Bay Area is somewhat surprising given California's public finance system, which limits the effect of local property values on school resources. This suggests that supply responsiveness may well be larger in other states. Moreover, it suggests that mechanisms at work more generally, such as changes in household monitoring in response to changes in local property values, may provide strong incentives for schools even in the absence of performance-based accountability reforms.

The rest of the paper is organized as follows: The next section motivates our empirical approach. In Section 3, we set out the demand model, and in Section 4, describe its estimation and address important identification issues. Section 5 describes the rich data used in the analysis, and Section 6 discusses our demand estimates, explaining how these are used to construct a slope measure for each school, based on meaningful variation in the data. Section 7 discusses the supply-side regressions that yield our main results, and Section 8 concludes.

2 SCHOOL CHOICE AND COMPETITION

In this section, we motivate more fully the use of the slope of a school's quality demand curve as our preferred competition measure. We then describe how we use this to shed light on the direct effect of competition on public school performance.

The slope of a school's quality demand curve measures the change in demand in response to a change in the school's quality. To see why this provides natural measure of competition, consider a stylized model of a local education market. The agents on the demand side comprise households with children; the supply side consists of teachers and school administrators. Household preferences are defined over consumption, housing services and school quality, and households choose where to live and where to send their children to school (there may be private school alternatives) based on quality and cost. Suppose school quality is unidimensional, measured by standardized tests, conditioning on school and family inputs. Quality is the output of a public school education production technology that converts student characteristics, school

---

[11] We also generate alternative slope measures based on a demand model that excludes commuting considerations, showing that such measures have essentially no power to explain test scores, once we

resources, teacher quality, and discretionary teacher and administrator effort into a performance measure.[12]

Public school objectives are potentially complex, not least because public schools tend to be heavily regulated. For simplicity, we will treat the school as a single effort-making body, and consider two polar cases. At one extreme, public schools could aim to maximize school quality, given resources; at the other, public schools could aim to maximize rents, which are increasing in school revenues (and thus enrollment) and decreasing in effort cost.

If schools were quality maximizers, then quality would be set independent of considerations about the effect that it would have on enrollment. School personnel would simply set effort at the maximal level consistent with their continued participation in the schooling sector. Significantly, under quality maximization, the school production function would be seen directly in the data, and the slope would have no effect on school performance.

In contrast, under rent seeking, a school would face the following trade-off: by raising quality (through higher effort), it would increase enrollment, and if funding were on a per-pupil basis, this would lead to an increase in revenues. At the same time, higher quality would require higher effort, which is costly. How the school would resolve this trade-off in making its optimal effort choice would depend in part on the size of the marginal benefit of higher effort. This, in turn, would depend on the response of demand to higher quality – the slope measure we focus on. Clearly, in a rent-seeking world, the slope of the demand curve *would* affect school production via discretionary effort choice: as the slope steepened, so the school would have greater incentive to raise quality so as to avoid a significant loss of enrollment.[13]

This discussion motivates our main estimating equation, which relates measured school performance to a set of relevant determinants: the characteristics of students, school inputs, teacher characteristics, and neighborhood controls, including our measure of local competition, the slope of the demand curve. Teacher and administrator effort is not observed by the researcher, but (depending on school objectives) might be influenced by market conditions. We adopt the following linear production function specification:

$$(1) \qquad T^m = X_c^{m\prime}\gamma_c + X_n^{m\prime}\gamma_n + \gamma_E E^m + \varepsilon^m$$

---

condition on the main set of controls.

[12] In practice, school personnel have other choice variables than effort, such as school resources.

[13] In practice, it is unlikely that schools are pure rent-seekers. However, to the extent that they have some discretion over quality setting and depart from pure quality maximization, so the slope of the quality demand curve will, to some degree, influence school quality setting.

where T^m is school m's average test score, X_c^m is a vector of the characteristics of the children that attend the school m, X_n^m represents a set of neighborhood controls, E^m is the slope of the quality demand curve, γ_c, γ_n and γ_E are coefficients to be estimated, and ε^m is a random error term. This is the equation we will take to the data, with interest focusing on the γ_E coefficient.

Beyond the stylized model, there are a variety of reasons why public schools might be responsive to the slope of the demand curve, consistent with a positive estimate for γ_E. A high slope might make homeowners more sensitive to school behavior, leading to better monitoring and more active political involvement in elections for school board and other local officials. In some school finance regimes, property values determine local property taxes and thus influence school revenue; thus schools would have an incentive in such settings to care about the effect of school quality on local property prices. Conversely, there are clear reasons why schools might not be responsive: teaching and administrative positions afford a good deal of job security; and political pressure might be ineffective if elections were for positions (such as mayor) where schooling is only one among many competing issues.

Construction of the Competition Measure

It is an empirical question just how strong these incentives to respond to competition are. To shed light on this issue, we estimate the key equation (1) above using a rich data set from the San Francisco Bay Area. For this purpose, we construct a set of school-specific slopes in which demand is measured in terms of local housing values,[14] and quality is measured using school average test scores, conditioning on school and student characteristics. The corresponding slope for a given school captures the change in local house prices as school quality changes.

The slope measures are not directly observed in the data. Rather, we estimate them in the following way: first, rich Census data are used to estimate a flexible demand system, taking careful account of endogeneity issues on the demand side of the market and making explicit the way that individual demands aggregate up to form a housing market equilibrium using an equilibrium model of the housing market. We then use these demand estimates in combination with the equilibrium model to perform a series of simple counterfactual experiments. In particular, for each of 720 elementary schools in the Bay Area, we use the equilibrium model to conduct a simple counterfactual simulation, raising its average test score by one standard deviation (78 test score points on a mean of 522) and calculating the new housing market equilibrium. This has the effect of increasing house values in the corresponding school

---

[14] Alternatively, demand could be measured based on student enrollment (or even enrollment of specific types of student). We intend to explore alternative measures in subsequent analyses.

attendance zone; the resulting predicted change in house values, given the change in school quality, is then used in the estimating various specifications of the regression shown in equation (1).

Unlike prior work, it is important to note that our approach provides estimates of school-specific competition measures, rather than MSA-level average choice indices. Constructing our slope measures from a single financing regime, rather than looking across MSAs, has the advantage that one would expect the incentives to respond to competition to vary with financing: our approach will better allow us to identify the direct effect of competition. We note that under California's local public finance system, the marginal dollar comes from the state rather than from local property taxes. This might be expected to provide weaker incentives to respond to competition, as schools are less able to take advantage of quality improvements directly. In turn, it is likely to provide lower bound on incentives to respond to competition, worth remembering when it comes to interpreting the economic significance of our results.

3 DEMAND

To measure household preferences, we now turn to a model of the residential location decision of households in the Bay Area. In developing such a model, our goal is to provide the simplest analytical tool that can account for (i) heterogeneity in both household characteristics and the attributes of houses/neighborhoods and (ii) the endogenous determination of housing prices and neighborhood sociodemographic compositions.

To this end, we use the equilibrium model of an urban housing market developed in Bayer, McMillan, and Rueben (2009). This equilibrium model consists of two key elements: the individual household residential location decision problem and a market-clearing condition. While maintaining a simple structure, the model allows households to have heterogeneous preferences defined over housing and neighborhood attributes in a flexible way; it also allows for housing prices and neighborhood sociodemographic compositions to be determined in equilibrium. In estimating the model, we are careful to account for the correlation that arises between unobserved housing and neighborhood attributes and both housing prices and neighborhood composition as a result of sorting.[15] Having estimated the model, we then use it to

---

[15] A long line of theoretical studies, including important papers by Epple, Filimon and Romer (EFR) (1984, 1993), Benabou (1993, 1996), Anas and Kim (1995), Anas (2002), Fernandez and Rogerson (1996, 1998), and Nechyba (1999, 2000) have developed and used models of sorting to analyze the way that interdependent individual decisions in the housing market aggregate up to determine the equilibrium structure of a metropolitan area. In recent years, a new line of empirical research has sought to take these models to the data. Epple and Sieg (1999) develop an estimator for the equilibrium sorting model of EFR, providing the first unified treatment of theory and empirics in the literature. In the same vein, Sieg *et al.*

construct measures of the slope of the demand curve faced by each school, explained more fully below.

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., crime, land use, and topography). As above, let T_h represent the quality of the local elementary school as measured by its average test score. 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:

$$(2) \qquad \underset{(h)}{Max} \quad V_h^i = \alpha_X^i X_h + \alpha_T^i T_h + \alpha_Z^i Z_h - \alpha_p^i p_h - \alpha_d^i d_h^i + \xi_h + \varepsilon_h^i .$$

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

Each household's valuation of choice characteristics is allowed to vary with its own characteristics, z^i, including education, income, race, employment status, and household composition. Specifically, each parameter associated with housing and neighborhood characteristics and price, α^i_j, for $j \in \{X, T, Z, d, p\}$, varies with a household's own characteristics according to:

(2004) use this approach to explore the general equilibrium impacts of air quality improvements in the Los Angeles Basin.

$$(3) \qquad \alpha_j^i = \alpha_{0j} + \sum_{k=1}^{K} \alpha_{rj} z_k^i ,$$

with equation (3) describing household i's preference for choice characteristic j.

This specification of the utility function gives rise to a horizontal model of sorting in which household preferences are defined distinctly over each choice characteristic, including both school quality and neighborhood sociodemographic characteristics.[16] 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.[17] The additional flexibility of the horizontal model is especially relevant for this paper as it is the magnitude of the heterogeneity in preferences 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 is typical for models in the broad sorting 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.[18]

---

[16] The horizontal specification also captures the geography of the urban housing market 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 metropolitan area, thereby increasing equilibrium housing prices in neighborhoods near employment centers.

[17] 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 that 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.

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

Given the household's problem described in equations (2)-(3), 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

$$(4) \qquad V_h^i > V_k^i \quad \Rightarrow \quad W_h^i + \varepsilon_h^i > W_k^i + \varepsilon_k^i \quad \Rightarrow \quad \varepsilon_h^i - \varepsilon_k^i > W_k^i - W_h^i \qquad \forall \quad k \neq h$$

where W^i_h includes all of the non-idiosyncratic components of the utility function V^i_h. As the inequalities in (4) imply, the probability that a household chooses any particular choice depends in general on the characteristics of the full set of possible house types. Thus the probability P^i_h that household i chooses housing type h can be written as a function of the full vectors of housing and neighborhood characteristics (both observed and unobserved) and prices $\{\mathbf{X}, \mathbf{Z}, \mathbf{p}, \xi\}$:[19]

$$(5) \qquad P_h^i = f_h(z^i, \mathbf{Z}, \mathbf{X}, \mathbf{p}, \xi)$$

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

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

$$(6) \qquad D_h = \sum_i P_h^i \ .$$

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

$$(7) \qquad D_h = S_h, \quad \forall h \quad \Rightarrow \quad \sum_i P_h^i = S_h \quad \forall h \, .$$

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

---

[19] 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.

fixed set of housing and neighborhood attributes, Bayer, McMillan, and Rueben (2009) show that a unique set of prices (up to a scale) clears the market.

Given that some neighborhood attributes are endogenously determined by the sorting process itself, we define a sorting equilibrium as a set of residential location decisions and a vector of housing prices such that the housing market clears and each household makes its optimal location decision given the location decisions of all other households. In equilibrium, the vector of neighborhood sociodemographic characteristics along with the corresponding vector of market clearing prices must give rise to choice probabilities in equation (4) that aggregate back up to the same vector of neighborhood sociodemographics.[20] Whether this model gives rise to multiple equilibria depends on the distributions of preferences and available housing choices as well as the utility parameters. In general, it is not possible to establish that the equilibrium is unique *a priori*. However, estimation of the model does not require the computation of an equilibrium nor uniqueness more generally, as we describe in the next section. We discuss the issue of uniqueness further in the context of calculating our demand slope measures below.

4 DEMAND ESTIMATION

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

$$V_h^i = \delta_h + \lambda_h^i + \varepsilon_h^i \tag{8}$$

where

$$\delta_h = \alpha_{0X} X_h + \alpha_{0Z} Z_h - \alpha_{0p} p_h + \theta_{bh} + \xi_h \tag{9}$$

and

$$\lambda_h^i = \left(\sum_{k=1}^{K} \alpha_{kX} z_k^i \right) X_h + \left(\sum_{k=1}^{K} \alpha_{kZ} z_k^i \right) Z_h - \left(\sum_{k=1}^{K} \alpha_{kp} z_k^i \right) p_h - \left(\sum_{k=1}^{K} \alpha_{kd} z_k^i \right) d_h . \tag{10}$$

---

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

In equation (9), δ_h captures the portion of utility provided by housing type *h* that is common to all households, and in (10), *k* indexes household characteristics. When the household characteristics included in the model are constructed to have mean zero, δ_h is the mean indirect utility provided by housing choice *h*. The unobservable component of δ_h, ξ_h, captures the portion of unobserved preferences for housing choice *h* that is correlated across households, while $\varepsilon_h{}^i$ represents unobserved preferences over and above this shared component.

The first step of the estimation procedure is equivalent to a Maximum Likelihood estimator applied to the individual location decisions, taking prices and neighborhood sociodemographic compositions as given.[21] This returns estimates of the heterogeneous parameters in λ and mean indirect utilities, δ_h. The estimator is based simply on maximizing the probability that the model correctly matches each household observed in the sample with its chosen house type. In particular, for any combination of the heterogeneous parameters in λ and mean indirect utilities, δ_h, the model predicts the probability that each household *i* chooses house type *h*. We assume that $\varepsilon_h{}^i$ is drawn from the extreme value distribution, in which case this probability can be written:

$$(11) \qquad P_h^i = \mathbf{exp}(\delta_h + \hat{\lambda}_h^i) \Big/ \sum_k \mathbf{exp}(\delta_k + \hat{\lambda}_k^i)$$

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

$$(12) \qquad \tilde{\ell} = \sum_i \sum_h I_h^i \, \mathbf{ln}(P_h^i)$$

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

---

[21] 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 $\boldsymbol{\varepsilon}$ is uncorrelated with the prices and

The Endogeneity of School Quality and Neighborhood Sociodemographic Composition. Having estimated the vector of mean indirect utilities in the first stage of the estimation procedure, the second stage involves decomposing δ into observable and unobservable components according to the regression equation (9).[22] In estimating equation (9), important endogeneity problems need to be confronted. To the extent that house prices partly capture house and neighborhood quality unobserved to the econometrician, so the price variable will be endogenous. Estimation via least squares will thus lead to price coefficients biased towards zero, producing misleading willingness-to-pay estimates for a whole range of choice characteristics. This issue arises in the context of any differentiated products demand estimation and we follow the approach described in Bayer *et al.* (2007) to instrument for price.

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

There are plausible reasons to think that households will sort with respect to such boundaries – an important theme in the work of Nechyba (1999, 2000), for example. Thus, while the boundary fixed effects are likely to control well for differences in unobserved fixed factors, neighborhood sociodemographics are likely to vary discontinuously at the boundary. This is important: it implies that boundary fixed effects isolate variation in both school quality and neighborhood sociodemographics in a small region in which unobserved fixed features (e.g., access to the transportation network) are likely to vary only slightly, thereby providing an appealing way to account for the correlation of both school quality *and* neighborhood sociodemographics with unobservable neighborhood quality.

We incorporate school attendance zone boundary fixed effects θ_{bh} when estimating equation (9). In particular, we create a series of indicator variables for each Census block that

---

neighborhood sociodemographic characteristics that arise in any equilibrium. For more discussion, see the Technical Appendix in Bayer *et al*. (2007).

[22] 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.

equal one if the block is within a given distance of each unique school attendance zone boundary in the metropolitan area. After first describing the basic features of the dataset, we show the variation in school quality and neighborhood sociodemographics at school attendance zone boundaries in the next section, drawing attention to clear discontinuities in neighborhood sociodemographics at the boundary.

5 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 about 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 several 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.

For this analysis, we use school attendance zone boundaries. 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 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 generally expect our analysis to understate 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,[23] number of rooms, number of bedrooms, type of structure, and the age of the building. We use these housing characteristics directly, and also construct neighborhood variables, such as neighborhood racial, education and income distributions, based on the households within the same Census block group (a Census region containing approximately 500 housing units). We merge additional data describing local conditions with each house record, constructing variables related to crime rates, land use, local schools, topography, and urban density. The list of the principal housing and neighborhood variables used in the analysis, along with means and standard deviations, is given in the first two columns of Table 1.[24]

School Assignment and School Quality. In order to implement the boundary approach, we gathered school attendance zone maps for as many elementary schools as possible in the Bay Area, for the period around the 1990 Census.[25] Our final attendance zone sample consists of 195 elementary schools – just under a third of the total number in the Bay Area. From this boundary sample, we excluded portions of boundaries coinciding with school district boundaries, city boundaries, or large roads, since they could potentially confound our identification strategy.

For Census blocks falling within these attendance zones, we followed a simple procedure to assign each block to a boundary. For each block, we calculated the perpendicular distance from the block population centroid to the nearest school attendance zone boundary. We then located the closest 'twin' Census block on the other side of that boundary. If a given block had a lower score than its twin, it was designated as being on the 'low' side of the boundary; otherwise it was designated as being on the 'high' side of the boundary. We restrict attention to boundaries for which we have Census data on both high and low sides.

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

---

[23] 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 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.

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

[25] School attendance zone maps are not provided or catalogued by the State of California. Therefore, we contacted all local school districts and schools in the study area individually and requested detailed maps for each school attendance zone within a given district around the period of analysis. Subsequently, these maps were digitized.

further subsample, consisting of houses assigned to Census blocks within 0.10 miles of the closest attendance zone boundary. Although the 0.10-mile subsample includes approximately half the number of observations, it provides a closer approximation to the ideal comparison of houses on the opposite sides of the same street, though in separate attendance zones.

Column 3 of Table 1 shows averages for the 0.20 miles subsample, and Column 3 of Table 1a presents analogous numbers for the 0.10 miles subsample. When comparing these to the full Bay Area sample (column 1), it is clear that prices, test scores, ownership, house size, average income and percentage white are slightly lower in the boundary subsamples. This is due in large part to the absence of San Francisco from our boundary samples, given that it does not have well-defined attendance zones.

While we have an exact assignment of Census blocks to school districts, and further, an exact assignment of blocks to school attendance zones for around a third of the schools in our sample, for the remainder, we employ the following approach for linking each Census block to a public 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.

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 the amount 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 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 0.2 miles of a school district boundary.[26] Comparing the first column to the third column of the table, the houses near school attendance zone boundaries are reasonably representative of those in the Bay Area as a whole.

The fourth and fifth columns report means for houses within 0.2 miles of a boundary, comparing houses on the high versus low average test score side of the each boundary; the seventh column reports t-tests for the difference in means. Comparing these differences reveals that houses on the high side cost \$18,700 more (on a mean of \$250,000) and are assigned to schools with test scores that are 74 points higher on average.[27] Moreover, houses on the high quality side of the boundary are much more likely to be inhabited by white households and households with more education and income. These types of across-boundary differences in sociodemographic composition are what one would expect if households sort on the basis of preferences for school quality. While far less significant, other housing characteristics do vary across the boundaries as well. Consequently, we expect the use of boundary fixed effects to 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.

6 DEMAND ESTIMATES

We noted in Section 3 that 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.

Mean Preferences[28]

Taking the latter first, Table 2 reports estimates of mean preferences for two specifications of equation (9), including neighborhood sociodemographic variables but without then with boundary fixed effects. We focus on results using the sample of houses within 0.20 miles of a boundary as they are more precise than the results using the sample within 0.10 miles.

The coefficient on the average test score in these specifications returns the average that households are willing to pay for a standard deviation increase in the average test score

---

[26] We experimented with a variety of distances and report the results for 0.2 miles rather than narrower bands, as these were more precise due to the larger sample size.

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

[28] The preference estimates in Tables 2 and 4 are drawn from Bayer *et al.* (2007).

conditional on the sociodemographic characteristics of the neighborhood, which are also indicative of the sociodemographic characteristics of the local school.

The estimated mean preferences for average test score are almost identical to the coefficients from the hedonic price regression (see Table 3). When boundary fixed effects are included in the analysis, the estimate mean MWTP for school quality is $19.7 per month compared with the estimated effect of $17.3 on housing prices in the analogous hedonic price regression reported in the second column of Table 3.

That the resulting MWTP for school average test scores is relatively small is 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 was 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 the test score 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 second column of Table 2 shows 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 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.

---

[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.

[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.

Heterogeneity in Preferences

Table 4 reports the implied estimates of the heterogeneity in MWTP for the average test score and neighborhood sociodemographic characteristics across households with different characteristics for our preferred specification, which includes both neighborhood sociodemographic characteristics and boundary fixed effects.[31]

The estimates of the heterogeneity in the MWTP for neighborhood sociodemographic characteristics reveal that while all households prefer to live in higher-income neighborhoods, *conditional on neighborhood income* households prefer to self-segregate on the basis of both race and education. In particular, the estimates imply that college-educated households are willing to pay $58 per month more than those without a college degree to live in a neighborhood that has 10 percent more college-educated households. When combined with the estimated mean MWTP of $10 per month reported in the first row, this estimate implies that households at each level of educational attainment prefer neighbors with like education levels: while college-educated households would pay an addition $32 per month to live in a neighborhood that had 10 percent more college-educated households, households without a college degree would actually need *compensating* to live in a neighborhood with 10 percent more college-educated neighbors, to the tune of $26 per month. Note that the preference for self-segregation on the basis of educational attainment is somewhat stronger for college-educated households.

Similarly, the heterogeneity estimates imply that blacks are willing to pay $98 more per month than whites to live in a neighborhood that has 10 percent more black versus white households. The mean MWTP for such an increase is -$10.5 per month, primarily reflecting the negative valuation of the white majority. Thus $98 is the difference between the *positive* MWTP of black households for this change and the *negative* MWTP of white households, indicating that households have strong self-segregating racial preferences.[32]

Focusing on heterogeneity in tastes for school quality, a household's willingness-to-pay increases with income, the presence of children, education, employment, and age. Blacks have a significantly lower willingness to pay for school quality relative to whites, although this may be related to unobservable factors such as the substantial degree of wealth inequality across race.

---

[31] The full heterogeneous choice model includes 135 interactions between nine household characteristics and fifteen housing and neighborhood characteristics. In Table 4, we only report MWTP for test scores and sociodemographics which correspond to the core of our analysis.

[32] It is also important to point out that these interactions pick up any direct preferences for living near others of the same race (e.g., a recent immigrant from China may want to interact with neighbors who also have immigrated from China) as well as any unobservable neighborhood or housing amenities valued more strongly by households of this group (e.g., recent immigrants from China may have similar tastes for shops, restaurants, and other neighborhood amenities).

The presence of children increases demand for school quality. That it does not increase demand by a greater amount may reflect the fact that the presence of children also raises the desired levels of other forms of consumption. The parameter estimates not presented in the table, for example, reveal that households with children have a much greater demand for larger houses.

As one might expect, increases in household income and education (which may proxy better for lifetime income) are associated with increased demand for better schools. They are also associated with higher demand for more educated and higher-income neighbors.

Constructing Measures of the Slope of the Demand Curve Faced by Each School

Given the estimates of the demand system, we now calculate a measure of the slope of the quality demand curve faced by each school. For each of 720 elementary schools in the Bay Area, we use the equilibrium model to conduct a simple counterfactual simulation, raising each school's average test score by one standard deviation (78 points on a mean of 522) and then calculating the new housing market equilibrium. In every case, this has the effect of increasing house values in the corresponding school attendance zone. The slope measure that we use in the subsequent analysis consists of the estimated change in average local property values per standard deviation increase in test score. Across the 720 elementary schools, the corresponding increase in house values following this increase ranges from \$1,815 to \$8,040 (or \$7-\$31 in monthly house prices), with a mean of \$5,180 (\$20) and a standard deviation of \$910 (\$3.50).[33]

Figure 1 shows the geographic distribution of the estimated slope measures across the Bay Area. In the figure, the area of the circle indicates the magnitude of the estimated slope. In general, the slope of a school's demand curve (as measured by the gradient of house values with respect to school quality) will be a function of two features of its environment: (i) the willingness of the households that it serves to pay for improvements in school quality, and (ii) the availability of close substitutes in geographically proximate neighborhoods. The former type of variation in

---

[33] The basic structure of the computation of the new equilibrium consists of a loop within a loop. Having lowered the test score of a given school by a standard deviation, 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. These prices can then be used directly to construct partial equilibrium slopes. For general equilibrium slopes, we take these new prices and the initial sociodemographic composition of each neighborhood and go on to calculate the probability that each household chooses each housing type. Aggregating these choices to the neighborhood level, we also compute the corresponding predicted sociodemographic composition of each neighborhood. We 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 represent the new equilibrium. The change in house prices then allows us to construct a general equilibrium slope measure, allowing households to be mobile.

the slope measure is problematic from the point of view of estimating the productive effect of competition in that a household's willingness-to-pay (WTP) for school quality is likely to be correlated with the performance of its children on standardized tests for reasons that have nothing to do with the school itself. The latter form of variation is less problematic in that a school and its corresponding neighborhood in any quality range can be located such that it has either many similar or dissimilar neighboring school catchment areas. It is this latter form of variation that we would like to exploit in our analysis.

Looking directly at the model of residential sorting estimated above, it is clear that the household sociodemographic characteristics included in the model, such as parental education, increase a household's estimated WTP for school quality and therefore increase the estimated slope of the demand curve for schools that serve these households. Thus in every specification of the analysis that follows, we include a complete set of controls for neighborhood averages of *all* the household sociodemographic and other housing and neighborhood measures included in our demand estimation. Without including this full set of controls, the competition measure would be mechanically correlated with the average school test score: because highly educated households are willing to pay more for school quality *and* select in to schools with higher test scores, the slope of the demand curves for these schools is mechanically steeper. When this full set of controls is included in estimating our main regression equations, however, this mechanical correlation is eliminated; doing so has the effect of restricting the effective variation in our main competition measure to variation associated with the availability of close substitutes in the local market.

Figure 2 shows the geographic distribution of the remaining variation in the estimated competition measures across the Bay Area once the full set of sociodemographic and other housing and neighborhood attributes included in the demand model have been conditioned out (i.e., they are residuals from a regression of our slope measure on this complete set of controls). For those familiar with the Bay Area, even a quick glance comparing Figure 1 and Figure 2 reveals that the remaining variation in the competition measure is not a simply a function of neighborhood socioeconomic conditions. For expositional ease, we work with a standardized version of this conditional measure throughout the remainder of our analysis. (Means and standard deviations for the school-related variables summarized for the 720 elementary schools used in the subsequent analysis are shown in Table 5.)

To demonstrate that the variation in these conditional slope measures is indeed related to the notion of the proximity of close substitutes, we construct a series of dissimilarity indices based on average house prices, neighborhood income, and neighborhood education. In particular, using the neighborhood attendance areas for the nearest 10 schools, we construct a measure of the average absolute difference between the measure of the school in question and each of these neighbors, weighting by the inverse of the distance. Thus a high measure for a given dissimilarity index indicates that a school is quite distinct from its geographic neighbors in the corresponding dimension (e.g., neighborhood educational attainment or racial composition).[34]

These dissimilarity indices provide a simple (albeit crude) measure of the availability of close substitute neighborhoods in the local market. Table 6 shows the correlation between these dissimilarity indices and our conditional slope measure (once the full set of sociodemographic, housing, and other neighborhood measures used in estimating the demand model have been conditioned out). In each case, the correlation is negative, indicating that schools that are differentiated from their neighbors in terms of sociodemographic and housing characteristics tend to have lower slope measures *ceteris paribus*. This suggests that the remaining variation in our slope measure is indeed picking up the type of variation that we would like to exploit in estimating equation (1).

As a further check, to demonstrate that our conditional slope measures are based on variation in the data driven by the geography of the Bay Area rather than the types of households in the local neighborhood, we conducted the following diagnostic test. First, we re-calculated an estimate of the slope of the demand curve using a model of residential sorting that did *not* include distance to work (i.e., a model that ignored geography). We then included these alternate slopes in average test score regressions analogous to those reported below in our main performance regression table, Table 7. Results from this exercise are shown in Table 7b. In each case, the coefficients on the demand measure are essentially zero and statistically insignificant. This diagnostic test gives us confidence that any mechanical correlation associated with the increased willingness of some households to pay for school quality (e.g., highly educated) is eliminated by the inclusion of a full set of controls for the variables used in the demand estimation. It also makes clear the importance of including geography in the demand estimation; it is the notion of the availability of close substitutes in geographic space that forms the basis for the remaining variation in our competition measures.[35]

---

[34] The point of this exercise is not to suggest that the set of the 10 closest neighboring schools necessarily contains a school's closest substitutes. High quality schools separated by large distances may be closer substitutes for one another than a high quality and low quality pair of neighboring schools. We use this measure primarily to show that our estimated slope measure is correlated with a measure of the availability of close substitutes, acknowledging that this constructed measure is not itself a perfect proxy.

[35] Finally, it is worth noting that by construction our competition measures are complex functions of observable household and neighborhood variables, including location. In this way, they are generally insulated against bias that might normally come about due to the correlation of unobserved ability and

7 SUPPLY-SIDE REGRESSIONS

We now turn to our main regression analysis. In particular, we report results for the specification shown in equation (1) using a variety of school characteristics (outputs and inputs) as the dependent variable and a series of six distinct sets of control variables for each school: housing and neighborhood characteristics, neighborhood sociodemographics, school sociodemographics, employment access and geographic variables, additional interactions among neighborhood sociodemographics, higher-order price and income terms, and land use variables. (The specific variables included in each set of controls are listed in Appendix Table 1.)

Table 7 reports results when the average 4th grade test score is used as the dependent variable. As mentioned above, we are using the conditional slope measure – the residual from a regression of the slope on a complete set of controls for the neighborhood average of all household sociodemographic, housing, and other neighborhood variables used in estimating the demand side of our mode. Thus, this amounts to including this complete set of controls in all specifications reported in Table 7.

The first column of Table 7 includes only these demand model variables as controls. The second column adds seventeen additional controls for 4th grade school sociodemographics. These are included to account for the fact that the sociodemographic characteristics of the students in the school (important in the production of the test score) might differ from the corresponding neighborhood average sociodemographics. The third column adds controls for five measures of the employment access by education level as well as four direct geographic measures. These controls are included to account for the possibility that the households who live within the core of the Bay Area may be systematically different from those who live in outlying regions.[36] A full set of parameter estimates for the specification shown in the third column are shown in Appendix Table 2. The fourth column adds controls for interactions between neighborhood race, education and income measures. This ensures that the remaining variation in the slope measure is not an artifact of non-linearities in these important household sociodemographics. Column five adds higher-order house price and income terms, and finally, column six adds four local land-use measures.

The estimates reported in the six columns of Table 7 reveal a consistent pattern of results with the conditional slope coefficient estimate falling in the range of 6.1 - 8.5 and the t-statistic ranging between 4.1 and 4.8. These coefficients are reported for a standardized conditional

---

unobserved demand on the part of parents; our competition measures are constructed so as to avoid including this portion of the variation in the measures.

competition measure and thus the interpretation is that a one standard deviation increase in the competitiveness of a school's local environment within the Bay Area is associated with a 6- to 8.5-point increase in the average test score of the school – in the range 0.08 - 0.11 standard deviations. This is a sizeable effect, indicating that a school's performance is indeed strongly linked to the competitiveness of its local environment.

We use a specification corresponding to the third column of Table 7 for the remaining analysis conducted in the paper. The first six columns of Table 8 report a series of such regressions using various school input measures as the dependent variable. The first three columns relate to teacher experience and reveal that schools facing a flatter slope of their demand curve actually have significantly fewer of the most productive teachers (those with experience between 5-9 years) and tend to have a higher number of teachers who are just beginning and nearing the end of their careers. The point estimates for the effect of the slope on the dependent variables considered in the next three columns, the pupil-teacher ratio and teacher education variables, are highly insignificant and very small in magnitude. Thus, in general, schools facing a higher competition measure appear to produce higher test scores without any significant advantages in terms of observable school inputs.

As discussed above, that school input decisions are not strongly related to the conditional slope in the Bay Area is not especially surprising given the restrictive financing regime in place in California. The fact that schools facing greater amounts of local competitive pressures do not attract observably better teachers provides some assurance that the remaining variation in our competition measure is not simply picking up unobserved student or neighborhood characteristics. Because teachers can sort across schools, we might be worried, for example, if a higher demand slope was correlated with the presence of more experienced or better-educated teachers. If anything, it appears that increased competition may push better teachers away from a school, perhaps because competition is associated with more intense monitoring.

The final column of Table 8 provides another way to evaluate the possibility that the conditional slope measure is correlated with unobserved student/household characteristics. In particular, column 7 reports results for a specification that uses the average amount of income from capital sources in the corresponding neighborhood (the best proxy for wealth available in the Census) as the dependent variable. Importantly, this measure was not used in estimating the demand side of the model and therefore serves as a check to see whether the remaining variation in our slope measure proxies for the type of unobserved household characteristic that might be expected to positively affect test scores. As the results reveal, the point estimate is actually

---

[36] It is also worth noting that the inclusion of county fixed effects also does not affect the results.

negative in this case (if imprecisely estimated, with a t-statistic of -0.5), indicating that higher competition schools actually serve households with lower levels of capital income *ceteris paribus*. This evidence provides a further indication that the remaining variation in our slope measure is not simply proxying for unobserved household characteristics.

In the final two tables, we explore the relationship between the dissimilarity indices described above and the average 4th grade test score. Again, it is this type of variation in our competition measure that we would ideally like to exploit in estimating the main specification shown in Tables 7 and 8, although it is again important to remember that these dissimilarity indices provide only a crude measure of the availability of close substitutes that our competition measures are capturing.

Table 9 reports the results of a series of specifications that relate the standardized school competition measure to various dissimilarity indices and additional control variables. Not surprisingly, given the negative correlations in Table 6, these measures are negatively related to the competition measure and are significant in most cases. When various combinations of these dissimilarity measures are included directly in the average 4th grade test score regression in Table 10, they collectively enter negatively (for the most part) and significantly. This lends further support to the view that the positive coefficient estimates reported in Table 7 are in fact reasonable. That is, schools increasingly isolated in product space (relative to their ten closest geographic neighbors) perform significantly worse.

8 CONCLUSION

Numerous studies have addressed the policy-relevant question of whether greater choice will lead to improvements in school quality, the typical focus in prior work being on the *overall* effect of increased choice on school performance. In the current paper, we began with the observation that this overall effect can be decomposed into a component measuring demand responsiveness (how increased choice affects school competition) and a second component measuring supply responsiveness (the way that increased competition affects school performance). By focusing on the overall effect of choice, the previous literature has provided little guidance as to the effect of competition itself on school performance. Moreover, existing research suggests that the responsiveness of demand to increased choice is weak, leaving open the possibility that supply responsiveness may be altogether more powerful.

We then presented a new approach for measuring the direct effect of competition on school performance – the strength of supply responsiveness. Central to this approach is the construction of our preferred measure of local competition – a measure of the slope of the quality demand curve faced by each school – which captures the change in demand each school faces in response to a change in that school's quality. We obtained this using a flexible demand model that includes a market-clearing condition, estimated using very rich Census data. This competition measure is then used in a regression framework that relates measured school performance to student, school and neighborhood controls.

Our results provide the first estimates in the literature of the direct effect of increased competition on public school performance. We find evidence of a marked increase in test scores in response to an increase in the slope of the demand curve: a one standard deviation increase in the competitiveness of a school's local market within the Bay Area is associated with a 0.1 standard deviation increase in school performance. At the same time, these achievement increases are accompanied by, if anything, slight reductions in important inputs, helping allay concerns about the role of omitted variables in these regressions and indicating that productivity increases with greater competition.

These findings are robust to inclusion of many types of control. We showed that our slope measure has no effect in predicting neighborhood wealth, lending support to the notion that our preferred competition measure is not simply picking up unobserved household characteristics. In contrast, it is correlated with similarity indices that describe the extent to which a school is isolated geographically and in terms of 'product' space: the slope measures steepen the less isolated a school becomes, as close substitutes become more readily available. As one might expect, these similarity measures also have a positive effect on school performance. In terms of heterogeneous effects, we find that school responsiveness to increased competition is greater in more educated communities, suggesting that educated parents may be better able to monitor school personnel as competition increases.

Overall, our evidence is consistent with strong supply responsiveness on the part of public schools. This is relevant to the broader school choice debate, suggesting that the performance benefits of policies that increase the availability of substitutes and so steepen school quality demand curves may be significant.

REFERENCES

Barrow, Lisa and Cecilia Rouse (2004), "Using Market Valuation to Assess Public School Spending," *Journal of Public Economics,* 88(9-10): 1747-69.

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

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

Belfield, Clive and Henry Levin (2002), "The Effects of Competition on Educational Outcomes: A Review of US Evidence," National Center for the Study of Privatization in Education Report, Columbia University.

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

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

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

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

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

Borland, M and R. Howson (1992), "Students' Academic Achievement and the Degree of Market Concentration in Education," *Economics of Education Review*, 11: 31-39.

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

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

Epple, Dennis, 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*, 24: 281-304.

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

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

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

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

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

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

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.

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

Figlio, David, and Joe Stone (1999) "School Choice and Student Performance: Are Private Schools Really Better?" *Research in Labor Economics*, 18: 115-40.

Hanushek, Eric and Steve Rivkin (2003) "Does Public School Competition Affect Teacher Quality," in *The Economics of School Choice*, edited by Caroline Hoxby, Chicago, IL: University of Chicago Press.

Hoxby, Caroline M. (1994) "Do Private Schools Provide Competition for Public Schools?" NBER Working Paper No. 4978.d

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

Jepsen, Christopher (1999) "The Effects of Private School Competition on Student Achievement," mimeo, Northwestern University.

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.

Millimet, Daniel and Vasudha Rangaprasad (2005) "Strategic Competition Amongst Public Schools," mimeo, Southern Methodist University.

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 (2006), "Good Principals or Good Peers? Parental Valuation of School Characteristics, Tiebout Equilibrium, and the Incentive Effects of Competition Among Jurisdictions," *American Economic Review*, 96(4): 1333-1350.

Sieg, Holger, V. Kerry Smith, H. Spencer Banzaf and Randall Walsh (2005), "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.

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