INTERNATIONAL ECONOMIC REVIEW
Vol. 45, No. 4, November 2004

ESTIMATING THE GENERAL EQUILIBRIUM BENEFITS OF LARGE CHANGES IN SPATIALLY DELINEATED PUBLIC GOODS*

BY HOLGER SIEG, V. KERRY SMITH, H. SPENCER BANZHAF, AND RANDY WALSH[1]

Carnegie Mellon University and NBER; North Carolina State University and Resources for the Future; Resources for the Future; University of Colorado, USA

The purpose of this article is to report a new approach for measuring the general equilibrium willingness to pay for large changes in spatially delineated public goods such as air quality. We estimate the parameters of a locational equilibrium model and compute equilibria for alternative scenarios characterizing the availability of public goods within a system of communities. Welfare measures take into consideration the adjustments of households in equilibrium to nonmarginal changes in public goods. The framework is used to analyze willingness to pay for reductions in ozone concentrations in Southern California between 1990 and 1995.

1. INTRODUCTION

Over the past 20 years, there has been a growing interest in evaluating the efficiency of public regulatory programs.[2] This article is motivated by fundamental problems of methods that are currently used in cost–benefit analyses. These methods were designed to consider relatively small projects that could be evaluated within a partial equilibrium framework. Current regulations, especially those associated with environmental policy, have economy-wide impacts. Several analysts have recognized the importance of general equilibrium adjustments for the costs attributed to environmental regulations. For example, Hazilla and Kopp (1990)

* Manuscript received March 2002; revised February 2003.

[1] We would like to thank an editor of the journal, two anonymous referees, Pat Bajari, Pat Bayer, Eli Berman, Steve Berry, Trudy Cameron, Steven Coate, Dennis Epple, Raquel Fernandez, Don Fullerton, Roger Gordon, Dan Hamermesh, James Heckman, Jim Hines, Matt Kahn, Lester Lave, Lars Nesheim, Ariel Pakes, Roni Shachar, and Rob Williams, as well as participants at numerous seminars and conferences. Thanks are due to Casey Schmeier and Jaren Pope for research assistance. This research was supported by the National Science Foundation NSF-SBR-98-08951, the U.S. Environmental Protection Agency R-828103. Sieg also acknowledges support from the Alfred P. Sloan Foundation. Please address correspondence to: Holger Sieg, Graduate School of Industrial Administration, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213-3890. Phone: 412 268-6564. Fax: 412 268-7357. E-mail: *holgers@andrew.cmu.edu*.

[2] The most recent realization of this activity in the United States is Section 624 of the Regulatory-Right-to-Know Act of 2001. This legislation requires the Office of Management and Budget to submit annually a report on the costs and benefits of Federal regulations, considering them in the aggregate and by agency program. In Europe, benefit cost analysis has not been as widely adopted as in the United States.

find that the price and income effects of environmental regulations for air and water pollution in 1990 implied social costs that were 2.6 times higher than the U.S. Environmental Protection Agency's estimates of annual compliance costs. More recently, Goulder et al. (1997) report that general equilibrium effects of preexisting taxes can increase the social costs of the control of SO_2 emissions by 71% over what they would be without these taxes. Neither of these studies considers the benefits associated with the regulations they evaluated.

The purpose of this article is to report a new approach for measuring the general equilibrium willingness to pay for large changes in spatially delineated public goods such as air quality. Households select one of a finite number of differentiated communities. Conditional on that choice, they select housing as a continuous, homogeneous good. Households' preferences are heterogeneous and satisfy single-crossing properties. Any equilibrium must, therefore, satisfy boundary indifference, stratification, and ascending bundles properties. These conditions for a locational equilibrium allow us to estimate the structural parameters of the model and compute equilibria under alternative scenarios characterizing public good provision.[3] Our study is the first empirical attempt to estimate the benefits for large changes in local public goods that consistently incorporate unobserved preference heterogeneity in estimation, computation of general equilibrium responses to policy changes, and, ultimately, welfare measurement.[4]

General equilibrium effects on benefits are typically ignored because estimates of nonmarket values of environmental amenities are derived from a partial equilibrium framework. In the case of air quality, the previous empirical literature has relied primarily on hedonic studies.[5] The first stage of a hedonic study consists of estimating the price function with public goods treated as site-specific attributes. Since optimal choices satisfy standard tangency conditions, we can estimate the marginal willingness to pay for public goods involved. Most studies only implement the first stage of the hedonic analysis. To estimate Hicksian willingness to pay it is necessary to identify and consistently estimate marginal willingness to pay functions in the second stage. Implementing the last part of the analysis is,

[3] This analysis is thus in the tradition of Harberger's (1962) pioneering work on general equilibrium in public economics. Our empirical analysis is also in the tradition of Oates (1969), who studied whether public spending and taxation were capitalized into land values. Most of the studies following Oates (1969) find evidence that capitalization of public spending, community specific amenities, and taxes is prevalent.

[4] More recently, Nesheim (2001) has estimated a hedonic model of school choice. Bayer et al. (2002) have adopted a framework developed by Berry, Levinsohn, and Pakes (BLP) to analyze racial sorting in the Bay Area. Walsh (2002) analyzes open space policies in Wake county using a similar approach as discussed in this article. Bajari and Kahn (2002) focus on housing markets and racial segregation. Ferreyra (2003) estimates a locational equilibrium model with fixed housing supply to study voucher programs.

[5] Following early efforts to estimate the incremental willingness to pay for ozone reductions in Los Angeles (Brookshire et al. 1982), the majority of the revealed preference estimates for the benefits from air quality changes use measures of particulate matter to characterize air quality conditions. For a review and a meta-analysis of these early studies see Smith and Huang (1995). Recent studies include Boyle et al. (1999), Chattopadhyay (1999), Chay and Greenstone (2000), and Beron et al. (2001).

however, problematic.[6] Up to now, there have been few successful applications. Our approach provides an alternative to hedonic methods.

Air quality improvements in Southern California between 1990 and 1995 were dramatic with ozone reductions ranging between 3 and 33% for 92 of the 93 school districts in the five counties in our sample. Our general equilibrium estimates of Hicksian willingness to pay for these ozone reductions are between 1% and 3% of annual household income. Average gains for the improvement of ozone conditions over these 5 years range from approximately \$120 to over \$9000 on the community level. The general equilibrium price adjustments predicted for these changes are substantial. The lowest price community experienced the largest increase in housing prices of 10.8% with its ozone concentrations declining by 24%. There were communities that experienced larger ozone improvements with smaller price increases. The largest was 33%, but this community had a small price increase because its price was at the highest level prior to the ozone improvement. These general equilibrium adjustments substantially change the distribution of benefits across households from partial equilibrium measures. We also find that poorer households can experience welfare losses because housing price increases offset air quality gains.

The remainder of the article is organized into seven sections. Section 2 discusses the basic framework used to describe household behavior and describes the estimation strategy. Section 3 discusses the construction of welfare measures. Section 4 describes the unique data set on housing prices and characteristics, air quality, and public education that is available for Southern California. Section 5 summarizes our estimation results. Section 6 outlines the general equilibrium benefit measurement. Section 7 discusses limitations of our framework and future research. Finally, Section 8 offers some concluding remarks.

2. ESTIMATING LOCATIONAL EQUILIBRIUM MODELS

2.1. *Communities and Household Preferences.* The model considers the problem of public good provision and residential decisions in a system of multiple jurisdictions.[7] The economy consists of a continuum of households living in a metropolitan area. The homogeneous land in the metropolitan area is divided among a finite number of communities, each of which has fixed boundaries. Jurisdictions may differ in the amount of land contained within their boundaries.[8] Households differ in their endowed income, y, and in a taste parameter, α, which

[6] These issues are discussed in Epple (1987), Bartik (1987), Palmquist (1991), and Ekeland et al. (2002).

[7] This literature was inspired by Tiebout (1956). See, for example, Epple et al. (1984), Goodspeed (1989), Epple and Romer (1991), Nechyba (1997a, 1997b) and Fernandez and Rogerson (1996, 1998). Locational equilibrium models were first estimated by Epple and Sieg (1999). Our approach builds on those contributions.

[8] We assume that the number of households is fixed in the metropolitan area. The model, thus, does not consider migration into the set of communities comprising the local market. To allow for migration into the metropolitan area, one could add an outside option along the lines suggested by Epple and Romano (1998).

reflects the household's valuation of the public good. The continuum of households is described by the joint distribution of y and α, which has continuous density, $f(\alpha, y)$. In our application, we assume that the joint distribution is bivariate log-normal. A household with taste parameter α and income y is referred to as (α, y). A household living in a community has preferences defined over a local public good, g, a local housing good, h, and a composite private good, b.

Denote with p the gross-of-tax price of a unit of housing services in a community. Households pay taxes that are levied on the value of housing services. Let t be an ad valorem tax on housing. Thus, $p = (1 + t)p^h$ with p^h the net-of-tax price of a unit of housing. The preferences of a household are represented by a utility function, $U(\alpha, g, h, b)$ that is twice differentiable in its arguments and strictly quasiconcave in g, h, and b. Households maximize their utility with respect to the budget constraint:

$$\max_{(h,b)} U(\alpha, g, h, b) \quad s.t.\ (1+t)p^h h = y - b \tag{1}$$

It is convenient to represent the preferences of a household using the indirect utility function derived by solving the optimization problem given in Equation (1). We assume that the indirect utility function of a household living in community j is given by the following parametric expression:

$$V(\alpha, y, g_j, p_j) = \left\{\alpha g_j^{\rho} + \left[e^{\frac{y^{1-\nu}-1}{1-\nu}} e^{-\frac{\beta p_j^{\eta+1}-1}{1+\eta}}\right]^{\rho}\right\}^{\frac{1}{\rho}} \tag{2}$$

This functional form allows for a general substitution pattern between private and public goods. The subutility function characterizing private consumption allows for constant income and price elasticities of housing, which are different from unity in absolute value.

Consider the slope of an "indirect indifference curve" in the (g_j, p_j)-plane:

$$M(\alpha, g_j, p_j, y) = \left.\frac{dp_j}{dg_j}\right|_{V=\bar{V}} \tag{3}$$

If $M(\cdot)$ is monotonic in y, then, for given α, indifference curves in the (g_j, p_j)-plane satisfy the "single-crossing" property. Similarly, the monotonicity of $M(\cdot)$ in α provides single-crossing for given y.[9]

In contrast, the Willig (1978) condition, an alternative restriction required to use weak complementarity to recover estimates of Hicksian consumer surplus for nonmarket resources, assumes that this same slope is invariant with income. The specification of the household decision problem and our indirect utility function in (2) implies that households consider public goods and housing prices at the extensive margin as a choice among a finite set of alternative communities. Households

[9] For a discussion of single-crossing properties, see Epple et al. (1984) or Epple and Romer (1991).

also choose the optimal amount of housing that is independent of the public good index, g_j, conditional on community choice.

In this article, we treat public good provision as predetermined.[10] As a result, we only analyze how households sort themselves into communities given the levels of public good provision and the resulting equilibrium in local housing markets. The single-crossing properties permit a structure in which necessary conditions for an equilibrium can easily be characterized. Let (g_i, p_i) and (g_j, p_j) be the level of public good provision and gross-of-tax housing price in communities i and j, respectively, and suppose that some households prefer (g_j, p_j) and others prefer (g_i, p_i). Then locational choices in equilibrium will satisfy three properties: boundary indifference, stratification, and ascending bundles, each described below.

2.2. *Community Size and Household Sorting.* If we order communities by ascending (g, p) pairs, the set of households indifferent between any two "adjacent" communities is given by the set of (y, α)s such that

$$V(g_j, p_j, y, \alpha) = V(g_{j+1}, p_{j+1}, y, \alpha) \tag{4}$$

For the parameterization of the indirect utility used in this study, the indifference locus satisfies

$$\ln(\alpha) - \rho\left(\frac{y^{1-\nu} - 1}{1-\nu}\right) - \ln(Q_{j+1} - Q_j) = K_j \tag{5}$$

where Q_j and K_j are given by

$$Q_j = e^{-\rho \frac{\beta p_j^{\eta+1} - 1}{1+\eta}} \tag{6}$$

$$K_j = -\ln\left(g_j^{\rho} - g_{j+1}^{\rho}\right) \tag{7}$$

Stratification implies that the population living in community j can be obtained by integrating between the two lines that characterize boundary indifference between adjacent communities. More formally, define the set of agents living in community j as

$$C_j = \left\{(\alpha, y) \mid K_{j-1} + \ln(Q_j - Q_{j-1}) \leq \ln(\alpha) - \rho\left(\frac{y^{1-\nu} - 1}{1-\nu}\right) \leq K_j + \ln(Q_{j+1} - Q_j)\right\}$$

The measure of households living in community j is therefore given by

$$P(C_j) = \int_{C_j} f(\alpha, y)\, d\alpha\, dy \tag{8}$$

[10] Epple et al. (2001) consider the case of endogenous public good provision via majority rule.

Similarly, we can characterize the income distribution in community j. The mean income level denoted by μ_j^y is given by

$$\mu_j^y = \int_{C_j} y f(\alpha, y)\, d\alpha\, dy / P(C_j) \tag{9}$$

Let $\bar{y}_j$ denote the empirical counterpart to μ_j^y. Let us also assume that average income is measured with error ϵ_j^y. We then have the following relationship between predicted and observed mean income:

$$\mu_j^y = \bar{y}_j + \epsilon_j^y \tag{10}$$

Similarly, we can characterize the higher moments of the income distribution of each community as well as its quantiles.

2.3. *Housing Demand.* Roy's identity implies that individual housing demand is given by

$$\ln h_j^d = \ln(\beta) + \eta \ln p_j + \nu \ln(y) \tag{11}$$

Averaging across households in each community we have

$$\begin{aligned} \ln H_j^d &= \int \ln h_j^d(\ln y) f_j(\ln(y)) d\ln(y) \\ &= \ln(\beta) + \eta \ln p_j + \nu \mu_j^{\ln(y)} \end{aligned} \tag{12}$$

where $f_j(\ln(y))$ denotes the marginal density of log-income in community j. $\mu_j^{\ln(y)}$ is the mean of log-income. Let us assume that aggregate housing demand is measured with error. Hence we have

$$\ln H_j^d = \ln(\beta) + \eta \ln(p)_j + \nu \mu_j^{\ln(y)} + \epsilon_j^d \tag{13}$$

Note that this equation resembles a traditional housing demand function.[11] The main difference is that $\mu_j^{\ln(y)}$ reflects locational choices of households and, thus, properly controls for self-selection.[12] Again, we can also characterize the higher

[11] In traditional housing models, the error is typically viewed as a demand shock that gives rise to a standard endogeneity problem of housing prices. We can follow Fernandez and Rogerson (1998) and reinterpret our model as a two-stage game in which households chose a location in the first stage and housing in the second stage. If households learn about their housing demand shocks after they have chosen a community, then households will sort in equilibrium as described above. We, can therefore, also interpret the error in the housing demand equation as an average demand shock.

[12] Local housing markets are modeled as in the earlier literature in urban economics. See, for example, Carliner (1973), Polinsky (1977), Polinsky and Ellwood (1979), and Hanushek and Quigley (1980).

moments of the housing demand distribution of each community as well as its quantiles.

2.4. *Public Good Provision.* We assume that public good provision can be written as a simple index of observed community-specific characteristics, x_j, and an error term, ϵ_j^g:

$$g_j = g\left(x_j, \epsilon_j^g\right) \tag{14}$$

Observed characteristics include school quality, protection from crime, and environmental quality. The key assumption is that the function, $g(\cdot)$, is the same for all households. This assumption allows us to consider multiple community-specific amenities while at the same time maintaining the hierarchical nature of the model. Since we are ultimately interested in estimation, we need to parameterize the index function. One way to proceed is to assume that $g(\cdot)$ is a linear function in characteristics

$$g\left(x_j, \epsilon_j^g\right) = x_j \gamma + \epsilon_j^g \tag{15}$$

Since the scaling of public good provision is arbitrary, we can normalize the coefficient of the first component of the index (e.g., school quality) to be equal to one assuming that the first characteristic is a good.

2.5. *Inversion and Implied Levels of Public Goods.* Recall that the size of the population in community j is given by Equation (8). This system of equations can be solved recursively to obtain the implied levels of public good provision as a function of the parameters of the model, θ, the community sizes, $P(C_1), \ldots, P(C_J)$, and housing prices, $(p_1, \ldots, p_J)$. As in Lemma 2 of Epple and Sieg (1999), we obtain the following result:

Given a value of public good provision in the lowest community, g_1, the levels of public good provision, $\{g_j\}_{j=2}^{J}$, implied by the observed community sizes are recursively defined by the following equations:

$$g_j = g_j(g_{j-1}, P(C_j), p_j, p_{j-1} \mid \theta) \quad j = 2, \ldots, J \tag{16}$$

The implied levels of public good provision for the set of communities are thus defined as the levels of public goods, which are consistent with the observed community sizes and observed housing prizes given the structure of our model.

In practice, this inversion can only be done numerically since the integrals characterizing community sizes cannot be computed analytically. This introduces a simulation error into the analysis.[13]

[13] This inversion is similar to the share inversion used in the estimation of dynamic models (Hotz and Miller 1993) and in differentiated product models (Berry 1994).

2.6. *Estimation.* In this article, we extend the estimation procedure developed in Epple and Sieg (1999) and provide a new estimator for the structural parameters which efficiently exploits all orthogonality conditions of the model simultaneously. Given the parameterization of the model, the parameter vector θ contains the parameters of the utility function as well as the parameters characterizing the distribution of income and tastes.[14]

Given the implied levels of public goods calculated using the inversion algorithm, we can evaluate the orthogonality conditions derived above without having to compute an equilibrium of the model. In this article, we use seven types of orthogonality conditions: three housing expenditure quantiles (25th, 50th, and 75th), three income quantiles (25th, 50th, and 75th), and one based on the implied levels of public good provision. Let $y_j(q)$ denote the qth quantile of the log-income distribution of community j. Similarly, let $r_j(q)$ denote the qth quantile of the log-housing expenditure distribution of community j. Stacking the orthogonality conditions, we obtain the following vector:

$$m_j(\theta) = \left\{ \begin{array}{l} g_j - x_j'\gamma \\ y_j(25) - y_j^N(25) \\ y_j(50) - y_j^N(50) \\ y_j(75) - y_j^N(75) \\ \ln r_j^N(25) - \ln(\beta) - (\eta+1)\ln p_j - \nu \ln(y_j(25)) \\ \ln r_j^N(50) - \ln(\beta) - (\eta+1)\ln p_j - \nu \ln(y_j(50)) \\ \ln r_j^N(75) - \ln(\beta) - (\eta+1)\ln p_j - \nu \ln(y_j(75)) \end{array} \right\} \tag{17}$$

Additional orthogonality conditions can be formed if there are instruments, w_j, that are uncorrelated with the errors of the model. Provided we have a sufficient number of orthogonality conditions, we can estimate the parameters of the model using a GMM estimator (Hansen 1982) that is defined as follows:

$$\hat{\theta} = \text{argmin}_{\theta\in\Theta} \left\{ \frac{1}{J} \sum_{j=1}^{J} w_j m_j(\theta) \right\}' V^{-1} \left\{ \frac{1}{J} \sum_{j=1}^{J} w_j m_j(\theta) \right\} \tag{18}$$

The optimal GMM estimator is obtained by setting V equal to the covariance matrix of moments. Broadly speaking, the asymptotic covariance matrix of this estimator has three components, which arise from three sources of errors.[15] First, community sizes, income distributions, and housing expenditure distributions are estimated based on a cross-sectional sample of size N (e.g., the U.S. Census). Second, there is a simulation error because inversion of community sizes is only feasible using numerical techniques. The simulation error depends on the number of simulations, S. Finally, there are aggregate errors because of unobserved

[14] The level of public good provision in the first community is treated as an incidental parameter.

[15] Berry et al. (2002) provide a detailed econometric analysis of these issues. In principle, one could also use bootstrap methods to compute standard errors.

community-specific attributes, shocks in aggregate housing demand, or measurement error. The relevant sample size is J. In all practical applications, S and N will be large relative to J. Thus the main source of sampling error is due to the finite sample of communities in the metropolitan area.

Instrumental variable estimation is necessary for a number of reasons. If we view the errors in the housing expenditure function as a demand shock, housing prices will be correlated with the error terms. In that case, we could use housing supply shifters to generate instruments for housing prices in the demand equation. If we view the error terms as measurement error in the dependent variable, there is no obvious need for instruments.

Endogeneity problems also arise in the regression of the implied level of g on observed characteristics. In contrast to the IO literature, we do not treat prices as product characteristics. Housing prices are not on the right-hand side of the regression of implied levels of public goods on observed amenities. Instead, our dependent variables, the implied levels of public goods, are a constructed variable that depends on the estimated housing prices, as well as the mean and the variance of the taste distribution. In that sense, this regression is more closely related to a Box–Cox model than a Berry, Levinsohn, and Pakes (BLP) type model.

Endogeneity problems arise if important amenities are observed by households, but not by the econometrician. More likely, our measures of community-specific characteristics are exhaustive, but not perfect. In principle, we can fairly accurately characterize the dimensions that households care about (school quality, crime, environmental quality, distance to employment centers, etc.) and we have controlled for most of these factors in our analysis. But as discussed in detail in Section 4 of this article, some of our measures are not ideal.

In our estimation, we use functions of the community rank as instruments in the regression of implied levels of public goods on observables. These instruments are primarily used to identify the mean and variance of the taste distribution. This identification assumption is justified if the perturbations of equilibrium that are due to unobserved characteristics of the communities are small in the sense that they affect the equilibrium sorting and housing prices, but do not affect the rank of a community in the system, i.e., if unobservables are of second-order importance. In the absence of a quasi-experimental research design, this may be as good as possible.[16]

2.7. *Discussion.* The demand side of locational equilibrium models is closely related to hedonic and other differentiated product models. In hedonic models, the choice space is continuous. Preferences and technologies are smooth.[17] Hence optimal choices can be characterized by standard first- and second-order conditions.

[16] A discussion of IV estimation in Box–Cox type models can be found in Amemiya (1985). See also Bayer (2001), Black (1999), and Chay and Greenstone (2000) for an alternative discussion of instruments.

[17] In our approach, local public goods are only available in a finite number of distinct bundles. Optimal community choices are therefore not characterized by tangency conditions, but a sequence of inequality constraints. We can thus view hedonic models as smooth approximations of discrete locational equilibrium models.

Estimation of hedonic models has been hampered by a number of well-known identification problems discussed, for example, in Epple (1987). More recently, Ekeland et al. (2002) and Heckman et al. (2002) have developed new approaches that resolve some of the identification problems. The key insight of these two papers is that the inherent nonlinearities of hedonic models are sufficient to identify the parameters of certain hedonic models. In our approach we use hedonic housing price regressions in a much more limited way. As discussed in Section 4 of this article, we construct interjurisdictional price and quantity indexes of housing using simple hedonic regressions, a procedure that is relatively robust with respect to misspecification and omitted variable problems.

The demand side of a locational equilibrium model also shares many similarities with demand models for differentiated products in industrial organization.[18] One prominent approach in industrial organization is due to Berry (1994) and Berry et al. (1995, 1999). Bayer (2001) has applied this framework to study locational sorting in Northern California. Bayer et al. (2002) estimate a model of racial sorting and peer effects using restricted-use microdata from the U.S. Census.

The basic idea of BLP is to compute demand functions by aggregating individual demands derived from random utility models (McFadden 1974). The key component of BLP is to invert the market share equations to recover the mean utility levels associated with each product. The parameters of the model are then estimated using an IV estimator that controls for the fact that prices and unobserved product characteristics are likely to be correlated. The analog of the share inversion of BLP is the inversion of the community size equations to compute the implied levels of public good provision in our model.

Although there are some similarities between BLP and our approach, there are also a number of important differences. Our approach does not rely on community (=product)-specific idiosyncratic shocks for each household. Adding these additive separable error terms to the indirect utility function is clearly helpful, since it induces smoothness on the underlying problem. Hence, mean utility levels can be computed efficiently by a contraction mapping algorithm, which allows researchers to estimate demand systems with many sources of horizontal taste heterogeneity in the characteristics space. In contrast, our approach allows for more limited patterns of substitution among communities. However, additive separable error terms also often undermine the welfare calculations as discussed in detail in Petrin (2002). Turning off the idiosyncratic errors comes at a significant cost. In particular, there is no longer a simple contraction mapping that can be used to do the share inversion. Instead, one has to solve a complicated system of nonlinear equations. One of the key insights of Epple and Sieg (1999) is that this system is still computationally tractable for a large number of hierarchical locational equilibrium models.

The dimension of the characteristics space is fixed in our approach, independent of how many communities are in the model. In random utility models such as BLP, the dimension of the product space increases linearly with the number of

[18] The pioneering early papers on characteristics models with horizontal product differentiation are due to Gorman (1980) and Lancaster (1966).

products. This property is problematic in welfare analysis. In that sense, our work is more closely related to Breshnahan's (1987) work on vertical product differentiation and more recent work by Berry and Pakes (2002) and Bajari and Benkard (2002).

Our model is a mixed-discrete continuous choice model. Households choose among a finite set of communities. Conditional on choosing a community, households also demand a continuous housing good. In that respect, our work is closely related to earlier work on discrete-continuous models by Hanemann (1984). In addition, housing supply is not fixed, but adjusts to price changes. Of course, this is not the only way of modeling housing. Treating housing as discrete and supplied in fixed amounts is the main alternative. These two assumptions greatly simplify the nature of the analysis. The resulting economy is an assignment model and existence of equilibrium can be established under fairly general conditions as discussed in detail in Nechyba (1997a). Both assumptions are, therefore, useful, but also restrictive. Our framework does not require either of them.

3. WELFARE ANALYSIS OF LARGE POLICY CHANGES

3.1. *Partial and General Equilibrium Welfare Measurement.* In this section, we discuss how to construct general equilibrium welfare measures of large policy changes using our model.[19] Consider the case in which we exogenously change the level of public good provision in each community from g_j to $\bar{g}_j$. In our application, the change in public good provision arises from improvements in air quality that are due to federal and state air pollution policies. The conventional partial equilibrium Hicksian willingness to pay, WTP_{PE}, for a change in public goods is defined in Equation (19):

$$V(\alpha, y - \mathrm{WTP}_{PE}, \bar{g}_j, p_j) = V(\alpha, y, g_j, p_j) \tag{19}$$

A general equilibrium evaluation of the effects of large changes in air quality across the set of communities allows households to adjust their community locations in response to these changes. Such an analysis implies that housing prices can change as well. An evaluation of the policy change should reflect the price adjustments stemming from any changes in community-specific public goods. Equation (20) defines the general equilibrium willingness to pay, WTP_{GE}.

$$V(\alpha, y - \mathrm{WTP}_{GE}, \bar{g}_k, \bar{p}_k) = V(\alpha, y, g_j, p_j) \tag{20}$$

where k (j) indexes the community chosen in the new (old) equilibrium. Since households may adjust their location, the subscripts for ($\bar{g}_k$, $\bar{p}_k$) need not match (g_j, p_j). To estimate the general equilibrium willingness to pay, we must solve for the new equilibrium distribution of households and the associated new price vector. Households remaining in community j can be expected to experience new

[19] For a discussion of general equilibrium welfare analysis in hedonic models see Scotchmer (1985, 1986).

conditions because the level of amenities and prices in j change from (g_j, p_j) to $(\bar{g}_j, \bar{p}_j)$. Households moving from j to a new community k experience changes in amenities and prices from (g_j, p_j) to $(\bar{g}_k, \bar{p}_k)$. To estimate the general equilibrium WTP, we must carefully account for these types of changes.

Once we have estimated WTP for each household in the economy, we can characterize the complete distribution of WTP. In some cases, it will be desirable to aggregate information and focus on a few select moments of this distribution. For example, the mean general equilibrium WTP for a given community can be obtained by integrating over the relevant space of household types:[20]

$$\text{WTP}^j_{GE} = \int_{C^o_j} \text{WTP}_{GE}(y, \alpha) f(\alpha, y)\, d\alpha\, dy / P\big(C^0_j\big) \tag{21}$$

where C^0_j is the baseline community designation. These welfare measures contrast with what is usually estimated using hedonic price functions. In a hedonic equilibrium, the marginal price estimated with a hedonic price equation for a public good, such as air quality, will equal the marginal willingness to pay (MWTP). We can also compute MWTP estimates based on our framework. At the initial income level for each household and the utility level associated with their initial location, the MWTP is defined in Equation (22):

$$\text{MWTP}_j = \frac{V_{g_j}}{V_y} = -\frac{\partial y}{\partial g_j} \tag{22}$$

For our specification of the indirect utility function, we obtain

$$-\frac{\partial y}{\partial g_j} = \frac{y^{\nu} \alpha g_j^{\rho-1}}{V^{\rho} - \alpha g_j^{\rho}} \tag{23}$$

Moreover, these measures can be averaged, comparably to the process-defining average WTP in Equation (21), for the households originating in community j. Of course, since housing prices have a finite set of values in our framework, the link between MWTP and "marginal" prices is not comparable to a hedonic model. Nevertheless, we can compare MWTP estimates based on our locational equilibrium model with those based on hedonic housing model. This comparison provides another benchmark to evaluate our estimates. One final aspect of the welfare measurement concerns the distinction between owners and renters.

3.2. *Owners Versus Renters.* Our definitions of WTP^j_{GE} to this point assume that all households are renters. It is straightforward to change the definition of the welfare measures to allow for ownership effects. In that case, we would get an

[20] To assure consistency between partial and general equilibrium measures we summarize both sets of estimates over the baseline set of households in each community and report the averages across households based on their baseline community location.

additional income term in the indirect utility function, which would reflect capital losses or capital gains depending on the sign of the housing price change.

Alternatively, we can evaluate changes in capital gains and measure the importance of the owner/renter distinction by computing the differences in rental payments, ΔR_j, in Equation (24):[21]

$$\Delta R_j = (\bar{p}_j - p_j)H_j \tag{24}$$

This approach is inherently simpler to implement. Moreover, we do not need to impose additional assumptions specifying how ownership rates differ among various income groups.

3.3. *Implementation.* To implement the welfare analysis, we need to solve the model and characterize the equilibrium allocation of households among communities for different distributions of public goods. To close the model, we assume that the housing stock is owned by absentee landlords and that the housing supply is given by

$$H_j^s = l_j p_j^\tau \tag{25}$$

where l_j is a community-specific constant, which reflects the differences in land endowments and other fixed factors, and τ is the constant supply elasticity.[22]

Broadly speaking, a locational equilibrium consists of a vector of housing prices and an allocation of public goods such that (a) every household consumes the optimal amount of housing, (b) each household lives in one community and no household wants to move to a different community, and (c) the housing market clears in every community.

Closed-form solutions of equilibria do not exist. Hence we rely on numerical algorithms. To compute the equilibrium prices, we must solve a system of nonlinear equations given by the J local housing market clearing conditions. Given the hierarchical structure of the model, we can start the process with an initial guess for the housing price in the first community. We then compute the housing prices for all other communities such that the first $J - 1$ housing markets clear. The housing price of the first community is adjusted and the sequence repeated until

[21] The rent change defined in (24) may seem to provide a simple means to "adjust" partial equilibrium WTP to correspond to our general equilibrium measure. This conjecture is incorrect. Using the definitions for WTP_{GE} and WTP_{PE} in terms of expenditure function, $(WTP_{GE} - WTP_{PE})$ is given by $y(\alpha, \bar{g}_j, p_j, V^0) - y(\alpha, \bar{g}_k, \bar{p}_k, V^0)$. This expenditure difference incorporates adjustment in the amount of housing consumed in response to the difference between p_j and $\bar{p}_k$. The expression for ΔR_j in Equation (24) holds the consumption of housing constant and thus does not fully reflect the adjustment represented in the difference between PE and GE measures of willingness to pay.

[22] In this article, we do not estimate housing supply functions. In computations for the general equilibrium, we set τ, the supply elasticity, equal to 0, 1, or 2 and calibrate the community-specific housing supply constant, l_j, such that the predicted housing demand evaluated at the estimated price indexes for housing in each community equals the housing supply in the baseline equilibrium. In the context of the demand for residential land, Walsh (2002) incorporates estimated land supply functions into a locational equilibrium framework.

TABLE 1
DESCRIPTIVE STATISTICS OF THE HOUSING SAMPLE

Variable	Los Angeles	Orange	Riverside	San Bernardino	Ventura
Number of observations	59,466	21,560	19,644	16,304	7851
Market value of house	245,325	262,142	143,450	151,845	253,684
Number of bathrooms	1.89	2.14	2.06	2.09	2.22
Number of bedrooms	3.01	3.31	3.25	3.26	3.46
Lot size (square feet)	7758	6761	10,118	8866	9364
Building size (square feet)	1537	1718	1614	1595	1800
Swimming pool (proportion)	0.16	0.16	0.12	0.13	0.14
Fireplaces (proportion)	0.53	0.26	0.83	0.79	0.79
Age	37.6	23.7	9.4	16.2	17.3
One mile of coast (proportion)	0.025	0.054	0.0	0.0	0.026
Three miles of coast (proportion)	0.073	0.153	0.0	0.0	0.146
Commute time (minutes)	26.7	26.2	32.8	29.8	26.3

the last housing market clears. Once we have computed the new equilibrium, we construct welfare measures as discussed in the previous section.

4. DATA

Our analysis focuses on the Los Angeles Metropolitan Area, which consists of the area west of the San Gabriel Mountains and includes parts of five counties: Los Angeles, Orange, Riverside, San Bernardino, and Ventura. We assume that the school district corresponds to the community a household selects in making its locational decisions. To implement the model we need household characteristics by school district, quality measures for local public education, data on housing markets, and air quality measures.

4.1. *Housing Markets.* A comprehensive database on housing markets in the LA metropolitan area was assembled based on housing transactions collected by a firm, Transamerica Intellitech, that sells access to these data to realtors and financial institutions. These data contain housing characteristics and transaction prices for virtually all housing transactions in Southern California between 1989 and 1991. Table 1 reports means of the main variables in the housing sample by counties for these years.

One potential drawback associated with using California data relates to Proposition 13, which in 1978 limited property taxes to 1% of assessed value and limited growth in assessments. It has been argued in the literature that Proposition 13 created a lock-in effect on home owners. A household faces a tax on mobility because property taxes are based on the market value at the time of the last sale. If the current market value exceeds the assessed value, the revaluation creates additional mobility costs. O'Sullivan et al.'s (1995) detailed quantitative analysis of these lock-in effects indicated that they are small. "For the average household, a 13% inflation rate will lengthen the average time between moves by only approximately 2 months" (p. 138).

Closely related to the lock-in effect are questions relating to turnover in housing markets. Our model is intended to represent the locational choices of all households in each community. As a result, it is important to gauge whether a sample of current transactions is representative of all household housing expenditures as described by the underlying housing stock for each of the school districts. The pattern of housing expenditures across communities is quite close between the 1990 U.S. Census and our estimates, with a correlation of 0.99. However, prices tend to be uniformly higher in the U.S. Census. Across our 93 school districts, prices are 6 to 12% higher in the census (interquartile range). In addition, the homes are much younger in our data set. Fifteen percent of our houses were built within 1 year, whereas in the 1990 census only 3% of homes were built in 1989–1990. The oversampling of new homes relative to what would be desired for measures of the overall housing stock is not surprising because newly built homes will automatically show up in housing transactions, whereas older homes will only show up when they turn over. According to the 1990 U.S. Census, 70–80% of all households change houses within 10 years. Given the scope of our housing data, we expect to capture most of these housing transactions in our sample. Approximately 20–30% of the houses in the U.S. Census have a housing tenure that is greater than 10 years. By construction, our sample only contains a fraction of these houses.

One approach to address differences between our transaction sample and the U.S. Census is to construct weights for the transaction sample. The basic idea is to assign large weights for the observations in the transaction sample that appear to be undersampled based on estimates from the 1990 U.S. Census and small weights to observations that are relatively oversampled. We construct these weights such that the transaction sample matches three attributes of the stock of all owner-occupied housing in the U.S. Census: location, the number of bedrooms, and the age of the house. Hence, the observations on the housing sales are weighted by the relative weight of their school district. Conditional on school district, they are weighted by the number of bedrooms (in five categories) and the age of the house (eight categories), assuming age and bedrooms are independent.[23] We estimate the price indexes and the quartiles for housing moments using both weighted and unweighted samples.[24]

4.2. *Housing Prices.* Housing price indices that control for observed differences in the quantity and quality of housing consumed within and across communities are an important aspect of the empirical analysis of households' locational and housing choices. Our framework assumes that local public goods, such as education and air quality, can be unbundled from the effects of structural and locational characteristics such as the size of the house, commuting times, and the distance to the coast. This approach is equivalent to the assumption that the housing structure can be built anywhere, but the inherent characteristics of a location cannot. It uses

[23] The correlation between bedrooms and age was not available.

[24] As discussed below, air pollution readings are assigned to each house based on its location in relation to air pollution monitors. The weights applied for housing expenditures were also applied in estimating the average pollution readings in each school district.

TABLE 2

DESCRIPTIVE STATISTICS FOR THE 93 SCHOOL DISTRICTS

Variable	Mean	SD	Minimum	Maximum
Community size	0.011	0.031	0.00077	0.311
Test score	2.097	0.346	1.489	3.330
Ozone	0.150	0.040	0.091	0.232
Ozone exceedances	40.47	34.13	1.000	105.00
PM_{10}	46.88	10.23	30.29	71.23
Crime	623	157	228	1153
Price (owner weight)	2.492	0.872	1.000	6.672
Price (unweighted)	2.300	0.892	1.000	7.010
Income 25	24,488	8377	12,067	52,505
Income 50	42,556	13,032	21,451	92,058
Income 75	67,003	21,771	34,505	169,030
Housing 25	17,716	10,762	5000	80,000
Housing 50	22,381	17,322	7500	148,000
Housing 75	28,142	21,417	9200	175,000

NOTES: These descriptive statistics are the averages summarized for the 93 school districts in our sample.

The values in the table are for the owner-weighted case. The distribution of housing expenditures is for the same case.

an assumed continuum of choices among the characteristics of housing to develop a price index for a homogeneous unit of housing in each community. The market value of a specific house located in a given community can be converted into the imputed rent using the approach outlined in Poterba (1992). We control for structural characteristics of the house, denoted by z_{jn}, such as size and age, as well as locational characteristics such as proximity to the coast and employment. Let u_{jn} denote the unobserved housing characteristics. We assume that quality-adjusted units of housing are given by $\ln(h_{jn}) = \delta' z_{jn} + u_{jn}$.

Rent, r_{jn}, measured for a quality-adjusted unit is by definition the product of the adjusted housing price, p_j, and the number of quality adjusted housing units, h_{jn}. Using our specification for h_{jn} in this definition and taking logarithms provides a well-defined regression model:

$$\ln(r_{jn}) = \ln(p_j) + \delta' z_{jn} + u_{jn} \tag{26}$$

We can estimate the model above using ordinary least squares with fixed effects for each community. Based on the estimated community-specific fixed effects we construct housing price indexes for each community in the sample. We estimate prices using both weighted and unweighted data. The results for both sets of price indexes are reported in Table 2 along with other features of the communities.[25] Our findings indicate that the relative prices differ by as much as 6.7 to 1 across

[25] In Sieg et al. (2002) we discuss the specific details of this approach and compare it to alternative approaches for measuring price indexes. These comparisons suggest a high level of consistency across alternative price indexes in the ranks of communities by prices.

communities with the owner-weighted case (7 to 1 for the unweighted). The majority of the housing prices differ by small amounts.

4.3. *Air Quality.* Data for observed concentrations of ozone and particulate matter less than 10 micrometers in diameter (PM_{10}) were obtained from the California Air Resources Board monitoring records. Southern California has one of the most extensive air quality monitoring systems in the world. In the five counties of interest, no fewer than 45 monitors were measuring ozone each year from 1987 to 1992, and, beginning in 1987, no fewer than 19 were measuring PM_{10}.

Two issues arise in using these data in our models of community and housing choice. First, there is not a one-to-one correspondence between air quality monitors and school districts. Thus, an interpolation problem must be addressed in associating school districts with air quality levels based on spatially discrete measures. Fortunately, the large number of monitors minimizes the compromises created by interpolation. Half of all houses are within 4.5 miles of at least one monitor, and 90% are within 9.9 miles. Air pollution is measured using a centered 3-year average of pollution readings for the nearest monitor to each house in each year of the average. Our temporal window is 1989–91. We investigate the effects of distance-weighed pollution measures and found no significant difference in the conclusions derived with these measures compared to this temporal average of the nearest monitor's readings. The community estimate is then computed by averaging using the Census owner weights and unweighted averages for the houses sold in each community.

The South Coast air basin does exhibit significantly different air quality conditions along coastal areas in comparison with eastern portions of the area near the mountains. Indeed, the state's tradeable permit system, Reclaim, acknowledges the importance of wind conditions that give rise to the superior coastal air quality by requiring two permits from coastal locations to match one at interior regions of the area. The widely dispersed monitoring system allows us to capture these differences. One concern that could be raised with measures of air quality along coastal areas is that they do not exclusively measure the improved air quality, but also access to the coast. This issue is difficult to incorporate in our locational model because homes in individual school districts can be heterogeneous in their distance from the coast. We incorporate coastal distance measures into the first stage price regressions to control for these effects.

Three measures of air pollution—ozone concentration, ozone exceedances, and particulate matter—are considered in evaluating the effects of air quality. Ozone is measured in parts per million (ppm) as the average of the top 30 1-hour daily maximum readings at a given monitor during a year. We also consider the observed exceedances of the 1-hour federal standard for ozone. Particulate matter was measured by the annual geometric mean (in micrograms per cubic meter for particulate matter of 10 μm or less in size). Both pollutants have been well documented to impact health status and have been found to influence housing prices in hedonic studies (Smith and Huang, 1995; Beron et al., 2001). In general, the effects of particulate matter are through impacts on increased mortality rates and effects on materials. These effects have not been shown to have any threshold, so annual mean for particulates is often used in epidemiological and economic analyzes of its effects. In the case of ozone, however, human health effects are more likely to be triggered at higher levels. The focus on maximum concentrations provides the rationale for considering the average of the 30 highest ozone readings or the number of days violating the federal standard (a count of exceedances).

4.4. *School Districts and School Quality.* Table 2 reports some descriptive statistics for the main variables characterizing 93 school districts in the sample. We find that school districts differ significantly in their sizes. Community size is measured as the proportion of the total population in all 93 districts. The smallest districts only contain a few thousand households. The largest school district in the sample is LA Unified with more than 1.4 million households, more than 30% of the total population in the LA metropolitan area. Since LA Unified is relatively large, we also estimate a version of the model in which we break up LA Unified into 11 smaller units. These units correspond to the school districts that were formed when the LA Unified was reorganized in April 2000.

The current school finance system in California was mainly shaped by two events: the 1971 decision of the state supreme court in *Serrano vs. Priest* and the approval of Proposition 13 in 1978 by voters in California. Before *Serrano vs. Priest* local school districts had fiscal autonomy. After this ruling the state imposed limits on spending and taxation of local districts and allocated aid in order to offset inequalities in local spending. The basic idea was to achieve convergence of expenditures per student by increasing the aid to poorer districts and capping the amount of expenditure growth in the richer districts. In addition, Proposition 13 limited the growth rates of expenditures over the last 20 years. As a result of these events, today most school districts in California have lower per capita expenditures per student than school districts elsewhere in the United States.

With small differences in educational expenditures, a better measure of school quality would be based on outcomes rather than state formula spending. There exist substantial differences in school quality among districts, measured by test scores. There are a number of comprehensive tests that have been designed and implemented by the California Department of Education. The primary purpose of these statewide testing programs is to monitor the performance of schools and provide some information to parents. We construct measures of performance from standardized test scores for each school district using the 1992–93 California Learning Assessment System Grade Level Performance Assessment test. Each student taking this exam is assessed at one of six performance levels (with 6 the highest level). In Table 2, we report average math scores of the school districts in our sample which range from 1.46 to 3.33.[26]

[26] Using output measures of school quality comes at a cost. At least part of the differences in test scores may reflect differences in peer quality as discussed in Epple et al. (2003). Since the main goal of this article is to analyze air quality improvements, we abstract from these complications and treat school quality as predetermined.

5. ESTIMATION RESULTS

We estimate the parameters of the indirect utility function and the underlying joint distribution of income and tastes for public goods by matching the predicted income distributions, the predicted distributions of housing expenditures, and the implied levels of public good provision to the data. We implement estimators that differ by the choice of the orthogonality conditions, the choice set (93 versus 103 communities), the price index measures (weighted versus unweighted), and the measures for local public goods and amenities.

Table 3 reports the parameter estimates and the estimated standard errors for several different model specifications. The results reported in column I use housing price estimates based on unweighted data. These results are obtained by matching housing expenditures and public goods. In column II we add orthogonality conditions that are based on quantiles of the income distribution to the estimation algorithm. As discussed in the previous section, using weighted data may be preferable to using unweighted data. In columns III and IV, we repeat this exercise using owner-weighted housing price estimates. Column V uses the same approach as column IV but adds crime as an additional measure of public good provision. In columns VI and VII, we estimate the model using the larger sample of 103 communities to investigate how decomposing the largest school district in our sample (LA Unified) affects our estimates.

In general, all parameters have the expected signs and are estimated reasonably accurately. The point estimates of the income elasticity of housing ν range from 0.73 to 0.84. These estimates are consistent with Polinsky's (1977) early summary of the income elasticity estimates for consistent micromodels. The estimates in his overview range from 0.75 to 0.90. Our estimates for the price elasticity of housing, η, range from -0.16 to -0.01. These results are not as close to the early estimates reported by Polinsky. Although this estimate is about one-fifth the average he reports, it can be expected to be sensitive to the procedures used to adjust prices to measure the demand for a homogeneous bundle of housing services.[27] The estimates Polinsky selected as best addressed this issue by using the Bureau of Labor Statistics' index of the annual cost of a standardized package of housing. This index does not have the same resolution as our individualized hedonic price index. Moreover, he reports estimates by Carliner (1973) and others on either side of his central measure with some as low as -0.10. Our price elasticity estimate is at the lower end of the range of estimates in absolute magnitude. However, these earlier studies have either used selected microsamples such as FHA data bases or aggregate data without the ability to adjust for the heterogeneity in housing.

[27] Based on simple Hicks–Allen relationships between price and income elasticities, the price elasticity estimate seems low in absolute magnitude in comparison with the income elasticity. Our smallest estimates arise with the owner-weighted price index using only the orthogonality conditions for housing expenditures. In this case, we could interpret the measured income responsiveness as reflecting changes in permanent and not current income. In this case, they would be consistent with a smaller share of permanent income associated with housing.

TABLE 3

PARAMETER ESTIMATES

	Unweighted Sample				Owner-Weighted Sample									
	93 communities				93 communities						103 communities			
	I Housing		II Housing & Income		III Housing		IV Housing & Income		V Housing & Income		VI Housing		VII Housing & Income	
	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
μ_y	10.768		10.765	0.038	10.728		10.782	0.066	10.888	0.051	10.741		10.829	0.035
σ_y	0.318	0.008	0.652	0.023	0.428	0.028	0.745	0.019	0.690	0.023	0.393	0.013	0.755	0.013
μ_α	−0.416	0.158	0.408	0.191	0.627	0.324	0.833	0.172	0.562	0.229	−0.295	0.146	0.760	0.199
σ_α	0.352	0.021	0.514	0.053	0.502	0.155	0.357	0.025	0.492	0.024	0.321	0.030	0.355	0.027
λ	−0.203	0.038	−0.219	0.058	−0.208	0.083	−0.197	0.051	−0.247	0.097	−0.291	0.022	−0.207	0.029
ν	0.841	0.013	0.762	0.009	0.749	0.022	0.734	0.011	0.752	0.012	0.836	0.006	0.729	0.014
η	−0.162	0.058	−0.116	0.043	−0.012	0.008	−0.039	0.015	−0.019	0.083	−0.151	0.079	−0.037	0.017
β	1.301	0.199	2.909	0.266	2.719	0.673	2.817	0.323	2.435	0.445	1.205	0.132	2.971	0.444
ρ	−0.079	0.006	−0.046	0.003	−0.039	0.005	−0.022	0.002	−0.030	0.002	−0.084	0.003	−0.023	0.002
g_0	1.202	0.122	1.308	0.152	0.934	0.210	0.780	0.137	0.859	0.152	1.342	0.060	0.813	0.115
γ_1	−1.906	0.676	−0.913	0.723	−2.627	0.826	−2.951	0.779	−1.974	0.831	−1.614	0.428	−2.923	0.674
γ_2	0.000		0.000		0.000		0.000		−0.001	0.003	0.000		0.000	

We can treat income as a latent variable in estimation and just match housing expenditures and public good provision as in columns I, III, and VI of Table 3. These estimators assume that housing consumption and locational choices are based on permanent income instead of current income. The identifying assumption that we make is that the mean of permanent income in the population equals the mean of current income. Alternatively, we can assume that housing and residential decisions are based on current income. In that case income is observed, and we add quantiles of the observed income distributions as additional orthogonality conditions in the estimation. The results are shown in columns II, IV, V, and VII of Table 3. Our estimates suggest that permanent income is less dispersed than current income. Our estimate for σ_y when treating permanent income as a latent variable is much lower than the observed standard deviation of current income. This suggests that a large part of the variation in income may be transitory and may not affect the demand for housing.

The point estimate for the correlation between income and tastes for local public goods is negative and ranges from -0.29 to -0.19. The single-crossing property implies that high-income households have a stronger demand for local public goods and amenities than low-income households. Our negative estimates of ρ confirm this hypothesis. However, we also observe a large amount of income and housing heterogeneity within communities. This feature of our data implies that income and tastes for public goods cannot be strongly positively correlated. If they were, the model would predict almost complete stratification by income across communities. This hypothesis is clearly rejected by the data. To capture the amount of income and housing heterogeneity within communities, we need weak negative correlation between income and tastes. This finding is not implausible since high-income households can more easily substitute private goods for public goods. For example, they hire tutors for their children or send them to private schools. Private security services can be used to reduce the reliance on local police. Avoiding air pollution primarily requires locational adjustment. Nonetheless, enhanced indoor facilities such as indoor pools or spacious rooms can reduce the time outside exposed to air pollution.

The parameter capturing the air pollution in the public good index, γ_1, has the expected negative sign. Higher air quality increases the level of local public good provision.[28] The coefficient of ozone is typically significantly different from zero at reasonable levels of confidence. Adding additional measures of public good provision does not significantly alter these findings. In column V, we add crime measures to the public good index (γ_2). We find that crime is insignificant in our models. However, this finding may be primarily due to a lack of good crime data at the local level.

[28] In previous research, we also estimated the model using PM_{10} and ozone exceedances as alternative measures for air quality. We also found that if we included multiple measures of air pollution, we continued to estimate a negative sign for ozone, but the sign of PM_{10} reverses and was positive. This finding suggests that here too, as in the case of hedonic models (Palmquist, 1991), it is hard to identify the separate effects of different air pollution measures in our sample. We also found that the effect of ozone is estimated with much greater precision than the effect of PM_{10}.

TABLE 4

COMPARATIVE EVALUATION OF BENEFIT MEASURES

County/ Area	No Supply Response I Δozone	II WTP_{PE}	III WTP_{GE}	IV MWTP	V $\Delta\, g$	VI $\Delta\, p$	VII WTP_{GE}/y	Supply Response VIII WTP^2_{GE}
Study area	0.193	1210	1371	44	0.047	0.002	0.021	1291
Los Angeles	0.208	1472	1556	50	0.041	−0.005	0.023	1510
Orange	0.180	901	1391	44	0.037	−0.009	0.022	1340
Riverside	0.207	834	372	25	0.093	0.058	0.012	443
S. Bernadino	0.163	738	367	24	0.085	0.049	0.012	433
Ventura	0.062	164	725	26	0.035	−0.012	0.018	715

NOTE: All monetary estimates are in 1990 dollars.

6. THE BENEFITS FOR OZONE CHANGES—1990 TO 1995

Between 1990 and 1995, Southern California experienced significant air quality improvements. Ozone concentrations were reduced by 18.9% for the study area as a whole. Moreover, there were considerable differences in the spatial distribution of these gains. Ozone changes ranged across communities from a 2.7% increase to a 33% decline. In Los Angeles County, the number of days exceeding the federal 1-hour ozone standard dropped by 27% from 120 to 88 days. Thus, this case offers an ideal application to evaluate the potential importance of general equilibrium adjustments in response to large changes in public goods. Both the magnitude of the change and its diversity would call into question any partial equilibrium measures of WTP.

The first six columns of Table 4 summarize the key features of the welfare estimates for the communities in each county of our sample. These results are developed using the owner-weighted model estimated using the income and housing expenditure orthogonality conditions as reported in Column IV of Table 3. Column I of Table 4 reports the average proportionate change in ozone across the school districts in each county expressed as an improvement. Column II reports the partial equilibrium welfare estimate for households for the ozone change with prices held at their original value in each school district. Column III reports the general equilibrium estimate for the ozone change. These estimates include the price change due to the adjustment of households in equilibrium and assume a supply elasticity of zero. Column IV provides the average value of the MWTP using the initial community's price and air quality conditions. Using the new community conditions after adjustment, the average value of the MWTP declines, as expected, from 2 to 13% from these baseline levels. Column V provides the proportionate change in the public good index experienced in equilibrium by those who originally lived in the school districts in each county. Column VI summarizes the corresponding average of the proportionate price change for the same groups. Column VII reports the ratio of general equilibrium WTP to income.

The average value of the partial equilibrium measure for the welfare gain, computed across all school districts, is \$1210. There is a considerable variation in gains experienced by residents of particular school districts. When these results are averaged at the county level, the average partial equilibrium estimates range from \$164 in Ventura County to \$1472 in LA County. Orange, Riverside, and San Bernadino counties have average gains that range between \$700 and \$901. In contrast, the general equilibrium estimates for the area as whole are \$1371. General equilibrium WTP estimates range from \$367 in San Bernadino county to \$1556 in LA County. Within each county there is also diversity in the WTP estimates by school district. In LA County, for example, the WTP_{GE} estimates range from \$486 in Compton Unified to over \$9000 in Beverly Hills. At the school district level, the ratio of general to partial equilibrium measures ranges from 0.28 to 8.81, with an average discrepancy of nearly 50%. We thus conclude that partial and general equilibrium estimates can differ substantially.

Partial and general equilibrium WTP estimates can also have different signs. In one school district, households experience an estimated average loss of −\$12 in general equilibrium. While ozone was reduced by 24% in this school district, housing prices increased by nearly 11%. This school district remained the lowest priced area for housing both before and after the improvement in ozone. Thus, there was no lower priced alternative where these households could relocate. Of course, even in this case the distribution of the general equilibrium WTP within this school district included some estimated gains. Nonetheless, with the estimated median general equilibrium WTP at −\$1.78, more than half of the households experienced losses. The partial equilibrium estimate for this school district provides a quite different perspective, implying average gains that are solely due to the improvement in ozone levels of \$430.

Partial equilibrium measures fail to account for household adjustments and the resulting price changes. Prices tend to increase in districts with large improvements and decrease in areas with small improvements. Interestingly, the households originally located in the one district experiencing a decline in air quality nonetheless experienced a general equilibrium gain on average of \$2485 since it is possible for them to adjust and improve both their air quality and housing costs. Housing prices in the school district declined, but not as much as in some other districts, −1.8% as compared to the largest example of a decline of −4.8%.

Our estimated marginal willingness to pay for air quality is generally consistent with the large literature evaluating air quality with hedonic price regressions. As noted in Table 4, our overall average marginal willingness to pay is \$44. This is the estimated value for a 1-unit change in our measure of air quality (mean of the top 30 ozone hours measured in parts per billion). Since hedonic studies have not used this particular measure of air pollution, instead using counts of exceedances of daily ozone standards or even measures of particulate matters instead of ozone, our estimate is not directly comparable to the literature. However, if we consider all these measures as different proxies of one underlying composite (smog) that is important to people, we can still compare them using a common proportionate change in each pollutant. Accordingly, we consider 1% of the 1990 average Los Angeles value for each of the measure used in different studies (i.e., top 30 hours, exceedances, or particulates) and adjust all values to 1990 dollars. With this normalization of units, our estimate for MWTP would be \$67. This result compares to a wide range of values in the literature that range between \$8 to \$181. Banzhaf (2002), for example, using the same housing data and looking at the same time period as this study, but measuring ozone by the number of exceedances within a hedonic price model, estimates a marginal price and increment in exceedances that imply an adjusted MWTP of only \$8. However, using discrete choice random utility models instead of hedonic regressions, he finds MWTP estimates with the same data between \$18 and \$25. In an independent hedonic study for Los Angeles, Beron et al. (2001) estimate a value of \$66 for our normalized change in ozone exceedances, and these results closely match our estimates. In an early study for San Fransisco, Brucato et al. (1990) estimate a marginal value to pay of \$36 for a normalized change in exceedances. Many other studies have used particulate matter instead of ozone to measure air pollution. A meta-analysis of some of these studies by Smith and Huang (1995) estimates a marginal value of \$181, when predicting air quality and income conditions in Los Angeles and when again normalizing units as described above for comparison. Thus, our overall average estimate for the MWTP generally falls at the low end of the range of existing estimates and is close to an estimate for one recent hedonic study that also uses ozone to measure air pollution.

Figure 1 provides a graphical illustration of the difference of the price and public good combinations that are being evaluated in partial and equilibrium analysis. It shows the community choice sets for the 93 school districts. The first choice set is the baseline equilibrium in 1990. It is illustrated using a dotted line with diamonds. The second choice set is the partial equilibrium choice set that assumes that prices do not change, but ozone readings are at 1995 levels. It is illustrated using a solid line with circles. The third choice set is the result of the new equilibrium in 1995. It is represented using a dashed line with triangles. The partial equilibrium choice set holds housing prices constant. At each price we can evaluate the extent of the improvement in air quality by the horizontal distance between the baseline equilibrium and the partial equilibrium choice set. Horizontal distances cannot be used to compare the baseline equilibrium and the 1995 general equilibrium choice sets because households within the area can relocate to other school districts.

Table 5 illustrates this adjustment by tracing this spatial reallocation of households from the Long Beach school district. It reports the five school districts that households who initially lived in Long Beach choose in the new equilibrium. For each of the five school districts, we report the new price and public good combination. We also report the share of the initial population of Long Beach that moves to the district as well as partial and general equilibrium WTP estimates. Ozone declined in Long Beach by over 9%. Nonetheless, only 3% of the households initially assigned to the school district, based on their income and tastes for public goods, remain in Long Beach. Prices decline from 2.438 to 2.362, but the new communities offer either lower price or improved public good options or both. The general equilibrium willingness to pay exceeds the partial equilibrium measure for all households.

We have also considered how sensitive results are to the treatment of supply responses in the housing market. The last column in Table 4 recomputes the general equilibrium willingness to pay measures using a supply elasticity of 2. The

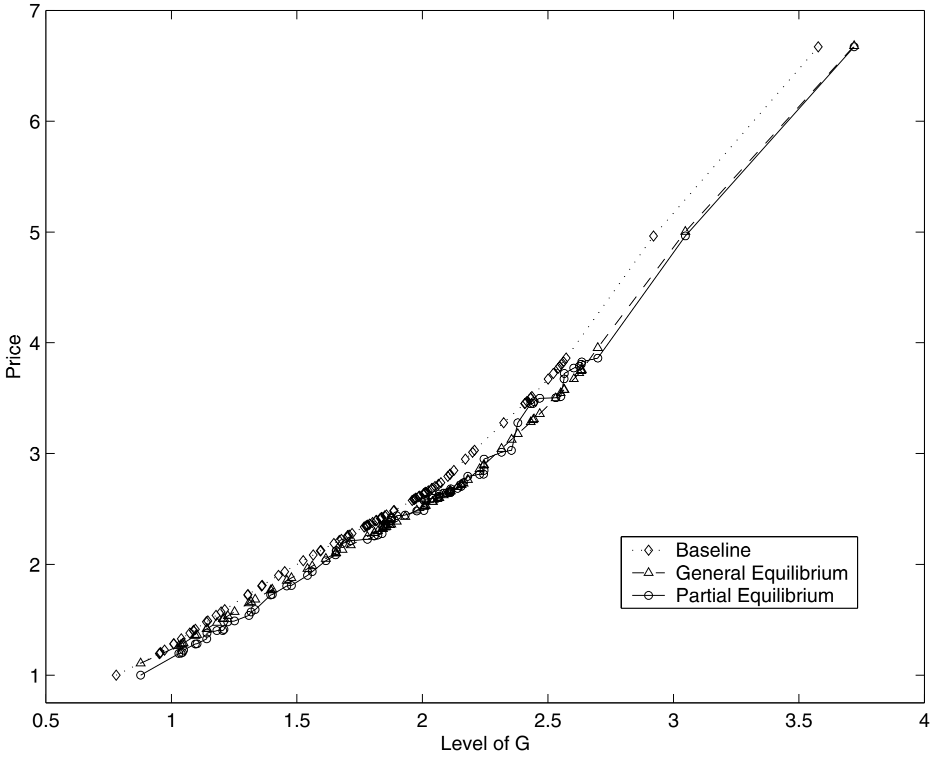

FIGURE 1

COMPARISON OF CHOICE SETS

differences among the two GE estimates are much smaller than those compared with the PE estimates. Comparing Column III with Column VIII, we find that the largest differences in the average WTP_{GE} estimates by county is \$71 dollars for Riverside County. For the other counties the differences are about \$10 to \$60. In contrast, the differences between Column II and Column III are over \$500. Hence we conclude that the average WTP_{GE} estimates are not sensitive to the choice of the housing supply elasticity.

TABLE 5

HOUSEHOLD RELOCATION AND WELFARE ANALYSIS: LONG BEACH

	p_2	g_2	Share of Population	WTP_{GE}	WTP_{PE}
Covina Valley	2.407	1.912	0.121	746	221
Oxnard Union High	2.393	1.910	0.424	746	220
Long Beach	2.362	1.876	0.034	719	208
Hacienda La Puerte	2.418	1.921	0.079	750	222
Whittier Union High	2.372	1.884	0.341	727	212

NOTE: Ozone reduction = 0.093; $p_1 = 2.438$ $g_1 = 1.851$.

TABLE 6
RENTS

	$\tau=0$	$\tau=1$	$\tau=2$
Study area	−102.19	−70.64	−70.39
Los Angeles	−74.40	−31.36	−24.86
Orange	−484.30	−444.18	−431.98
Riverside	472.37	445.87	402.68
San Bernadino	376.49	352.58	312.02
Ventura	−560.52	−547.00	−551.47

As Palmquist (1988) noted, changes in rents affect property owners differently than renters and need to be included in a comprehensive analysis of the benefits of policy changes. For the area as a whole, the average rent changes are small, −\$102.19 with zero supply elasticity and −\$70.39 with $\tau=2$. However, the distribution of these rent changes is quite pronounced. Table 6 summarizes this aspect of our model's implications. Average rent changes, which are defined as the expenditures for the initial amount of housing, vary substantially across counties, with Los Angeles experiencing the smallest changes in absolute magnitude and Ventura the largest.

Finally, we provide some additional sensitivity analyses. Table 7 reports the same welfare analysis summarized at the county level and for the area as a whole for the model based on price indexes from a hedonic model that does not reweight the data to match the Census and with an expanded 103 community choice set. The expanded choice set has little effect on the welfare measures. Price changes are comparable and both partial and general equilibrium welfare measures are not markedly impacted. The use of price indexes weighted to match the Census

TABLE 7
EFFECTS OF SAMPLE WEIGHTS AND CHOICE SETS ON WELFARE

County/Area	Δg	Δp	WTP_{GE}	WTP_{PE}	Δ Rents
			Unweighted Prices (93)		
Study area	0.012	0.003	665	618	−34
Los Angeles	0.011	0.000	806	781	−6
Orange	0.008	−0.008	661	440	−220
Riverside	0.020	0.016	148	313	167
San Bernadino	0.018	0.012	165	292	127
Ventura	0.007	−0.006	312	69	−240
			Owner-Weighted Prices (103)		
Study area	0.047	0.003	1376	1226	−141
Los Angeles	0.042	−0.004	1605	1471	−121
Orange	0.037	−0.010	1497	993	−503
Riverside	0.096	0.066	341	812	481
San Bernadino	0.081	0.048	351	731	386
Ventura	0.041	−0.007	918	209	−698

has a larger impact, reducing both general and partial equilibrium welfare measures. This is primarily due to the greater proportion of newer, smaller houses in the transaction sample than is present in the area's housing stock. Because the locational equilibrium model is intended to represent all housing choices, not simply those associated with new homes, this distinction is important to welfare analysis with the model. We view the results with price indexes estimated using Census-weighted data as the preferred measures.

7. LIMITATIONS AND FUTURE RESEARCH

Locational equilibrium models offer a parsimonious approach for estimating the distribution for a general equilibrium measure of household willingness to pay associated with large policy changes. These types of welfare measures necessarily require structural models consistent with restrictions implied by budget-constrained utility maximizations. Inevitably, meeting this goal introduces restrictive assumptions. These maintained hypotheses serve as both limitations of the current framework and opportunities for future research.

One important limitation arises from the restrictions on substitution among communities imposed by our preference specification. It seems reasonable to expect that households would differ in their relative preference for separate community-specific public goods, such as education, air quality, safety, etc. Our specification precludes these differences. We assume all public goods are represented with a single index and preference heterogeneity can be captured using a single parameter. This vertical differentiation facilitates estimation and computation of equilibria. Demand models allowing for horizontal differentiation, such as Berry et al. (1995), permit more general substitution and are important next steps in this type of research.

Parsimony necessarily raises the possibility of omission of important community-specific attributes. Peer group effects may well be the most important such omitted source of social interaction.[29] We estimated correlations between the community-specific residuals from the public good equation (i.e., an estimate measuring the unobserved component of local public goods) and several sociodemographic variables based on aggregates of the block-group information from the 1990 Census to the school district level. These included the proportion of married couples with children ages 3 to 11, as well as single-parent households with children, the proportion black, and two education measures—proportion without a high school degree and proportion with at least college. In general, the correlations between these peer measures and the residuals of the public good equation are small in absolute magnitude. The three largest correlations are for the two educational attainment measures estimated at 0.15 and the proportion of married families with children 3 to 11 measured at −0.21. These estimates indicate that peer effects may be a consideration, but at the school district level do not seem

[29] For a discussion of identification of peer effects see Manski (1993) and Brock and Durlauf (2001). Bayer et al. (2002) estimate a residential choice model with social interactions using microdata from the Bay Area.

to offer an important factor in defining attributes in addition to education and air quality that distinguished our communities.

We would also like to control for other spatially delineated factors, such as geographic and topographical features of the metro area, local labor markets, and transportation systems. Although each of these factors is potentially important, their omission is not unique to our model. Finally, households face mobility costs, which may prevent housing markets from clearing easily. If there are significant mobility costs, a smaller fraction of households will move in response to a large change in public goods. Our analysis has focused on two polar cases. The partial equilibrium measures correspond to the case in which moving costs are prohibitively high. The general equilibrium measures assume that moving costs are negligible. The PE and GE measures, therefore, can be interpreted as bounding the welfare measures that would correspond to a more realistic case with moving costs.

8. CONCLUSION

Benefits estimation is an increasingly important component of evaluating environmental and other types of regulations. In the United States, the Office of Management and Budget (OMB) uses benefit–cost analysis to evaluate regulations with costs exceeding \$100 million. OMB's most recent report on the costs and benefits of major rules assembled information from 67 regulations between 1995 and 2001. The majority of the aggregate costs estimated for these regulations relate to environmental rules. The sheer scope of these regulations, which involve about \$40 billion of annual costs, suggests that partial equilibrium analysis is not appropriate. Environmental policies usually involve large, spatially differentiated changes in the amounts of the public goods available. Households will undoubtedly adjust to these new conditions. Thus, the policies to be evaluated necessarily involve changes in prices as a result of changes in quantities of public goods. Locational equilibrium models offer a method to consistently incorporate unobserved heterogeneity in household preference in all three steps of the analysis required to develop general equilibrium benefit estimates. That is, the method provides estimates of heterogeneity in household preferences, uses these estimates in computing the general equilibrium consequences of a large policy change, and provides a consistent basis for incorporating preference heterogeneity in measures of the distribution of benefits.

Our empirical results suggest that there are large differences between partial and general equilibrium measures of willingness to pay for spatially differentiated improvements in air quality. These differences arise because households are mobile and adjust their housing demands and locational choices in response to a large policy change, with more wealthy households generally moving down to communities with previously low levels of public goods. This environmental gentrification leads to significant price increases in communities with large improvements in air quality and price decreases in communities with small air quality improvements. Distributional effects of environmental policies seem to be pronounced with opportunities for the lowest income households to lose because the induced increases in their housing prices are not fully offset by the air quality improvements they can afford to enjoy.

Broadly speaking, our general equilibrium analysis shows that renters in the low amenity and low income communities that have experienced the largest improvements in air quality are not necessarily the ones who gain the most from the policy changes. Their partial equilibrium gains due to the increase in public good provision can be offset by housing price increases. Property owners in the low amenity communities, however, gain significantly due to the appreciation of their property. This finding illustrates an important, but overlooked, dimension of the distributional analysis of environmental policies: It is not the residents of the low income communities who derive large gains from air quality improvements since they are renters. Instead the property owners, who may not even reside in the disadvantaged area, may well be the main beneficiaries of these policies.

REFERENCES

AMEMIYA, T., *Advanced Econometrics* (Cambridge, MA: Harvard University Press, 1985).

BAJARI, P., AND L. BENKARD, "Demand Estimation With Heterogenous Consumers and Unobserved Product Characteristics: A Hedonic Approach," Working Paper, Stanford University, 2002.

——, AND M. KAHN, "Estimating Housing Demand with an Application to Explaining Racial Segregation in Cities," Working Paper, Stanford University, 2002.

BANZHAF, H. S., "Quality Adjustment for Spatially-Delineated Public Goods: Theory and Application to Cost-of-Living Indices in LA," RFF Working Paper 02-10, 2002.

BARTIK, T., "The Estimation of Demand Parameters in Hedonic Price Models," *Journal of Political Economy* 95 (1987), 81–88.

BAYER, P., "Exploring Differences in the Demand for School Quality: An Empirical Analysis of School Choice in California," Working Paper, Yale University, 2001.

——, R. MCMILLAN, AND K. REUBEN, "The Causes and Consequences of Residential Segregation: An Equilibrium Analysis of Neighborhood Sorting," Working Paper, Yale University, 2002.

BERON, K., J. MURDOCH, AND M. THAYER, "The Benefits of Visibility Improvement: New Evidence from the Los Angeles Metropolitan Area," *Journal of Real Estate Finance and Economics* 22(2/3) (2001), 319–38.

BERRY, S., "Estimating Discrete-Choice Models of Product Differentiation," *Rand Journal of Economics* 25(2) (1994), 242–62.

——, J. LEVINSOHN, AND A. PAKES, "Automobile Prices in Market Equilibrium," *Econometrica* 63(4) (1995), 841–90.

——, ——, AND ——, "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy," *American Economic Review* 89 (1999), 400–30.

——, O. LINTON, AND A. PAKES, "Limit Theorems for Estimating Parameters of Differentiated Product Demand Systems," Working Paper, Yale University, 2002.

——, AND A. PAKES, "The Pure Characteristics Demand Model," Working Paper, Yale University, 2002.

BLACK, S., "Do Better Schools Matter? Parental Valuation of Elementary School Education," *Quarterly Journal of Economics* 114(2) (1999), 577–600.

BOYLE, K. J., P. J. POOR, AND L. O. TAYLOR, "Estimating the Demand for Protecting Freewater Lakes from Eutrophication," *American Journal of Agricultural Economics* 81(5) (1999), 1118–22.

BRESHNAHAN, T., "Competition and Collusion in the American Auto Industry: The 1955 Price War," *Journal of Industrial Economics* 35 (1987), 457–82.

Brock, W., and S. Durlauf, "Discrete Choice with Social Interactions," *Review of Economic Studies* 68(2) (2001), 235–60.
Brookshire, D., M. Thayer, W. Schulze, and R. d'Arge, "Valuing Public Goods: A Comparison of Survey and Hedonic Approaches," *American Economic Review* 72 (1982), 165–78.
Brucato, P., J. Murdoch, and M. Thayer, "Urban Air Quality Improvements: A Comparison of Aggregate Health and Welfare Benefits to Hedonic Price Differentials," *Journal of Environmental Management* 30 (1990), 265–79.
Carliner, G., "Income Elasticity of Housing Demand," *Review of Economics and Statistics* 55 (1973), 528–32.
Chattopadhyay, S., "Estimating the Demand for Air Quality: New Evidence Based on the Chicago Housing Market," *Land Economics* 75(1) (1999), 22–38.
Chay, K., and M. Greenstone, "Does Air Quality Matter? Evidence from Housing Markets," Working Paper, University of Chicago, 2000.
Ekeland, I., J. Heckman, and L. Nesheim, "Identification and Estimation of Hedonic Models," Working Paper, University of Chicago, 2002.
Epple, D., "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products," *Journal of Political Economy* 95 (1987), 59–80.
——, R. Filimon, and T. Romer, "Equilibrium Among Local Jurisdictions: Towards an Integrated Approach of Voting and Residential Choice," *Journal of Public Economics* 24 (1984), 281–304.
——, and R. Romano, "Competition Between Private and Public Schools, Vouchers and Peer Group Effects," *American Economic Review* 88 (1998), 33–63.
——, ——, and H. Sieg, "Admission, Tuition, and Financial Aid Policies in the Market for Higher Education in Equilibrium," GSIA Working Paper, 2003.
——, and T. Romer, "Mobility and Redistribution," *Journal of Political Economy* 99(4) (1991), 828–58.
——, T. Romer, and H. Sieg, "Interjurisdictional Sorting and Majority Rule: An Empirical Analysis," *Econometrica* 69 (2001), 1437–65.
——, and H. Sieg, "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy* 107(4) (1999), 645–81.
Fernandez, R., and R. Rogerson, "Income Distribution, Communities, and the Quality of Public Education," *Quarterly Journal of Economics* 111(1) (1996), 135–64.
——, and ——, "Public Education and Income Distribution: A Dynamic Quantitative Evaluation of Education-Finance Reform," *American Economic Review* 88(4) (1998), 813–33.
Ferreyra, M., "Estimating the Effects of Private School Vouchers in Multi-District Economies," GSIA Working Paper, 2003.
Goodspeed, T., "A Reexamination of the Use of Ability-to-Pay Taxes by Local Governments," *Journal of Public Economics* 38(3) (1989), 319–42.
Gorman, W., "A Possible Procedure for Analyzing Quality Differentials in the Egg-Market," *Review of Economic Studies* 47 (1980), 843–56.
Goulder, L. H., J. Parry, and D. Burtraw, "Revenue-Raising Versus Other Approaches to Environmental Protection: The Critical Significance of Pre-existing Tax Distortions," *Rand Journal of Economics* 28 (1997), 708–31.
Hanemann, M., "Discrete Continuous Models of Consumer Demand," *Econometrica* 52(2) (1984), 541–61.
Hansen, L. P., "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica* 50(4) (1982), 1029–53.
Hanushek, E., and J. Quigley, "What Is the Price Elasticity of Housing Demand?" *Review of Economics and Statistics* 62(3) (1980), 449–54.
Harberger, A., "The Incidence of the Corporation Income Tax," *Journal of Political Economy* 70 (1962), 215–40.
Hazilla, M., and R. J. Kopp, "The Social Cost of Environmental Quality Regulation: A General Equilibrium Analysis," *Journal of Political Economy* 98(4) (1990), 853–73.

HECKMAN, J., R. MATZKIN, AND L. NESHEIM, "Nonparametric Estimation of Nonadditive Hedonic Models," Working Paper, University of Chicago, 2002.

HOTZ, J., AND R. MILLER, "Conditional Choice Probabilities and Estimation of Dynamic Models," *Review of Economic Studies* 60 (1993), 497–529.

LANCASTER, K., "A New Approach to Consumer Theory," *Journal of Political Economy* 74 (1966), 132–57.

MANSKI, C., "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies* 60 (1993), 531–42.

MCFADDEN, D., "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics* (New York: Academic Press, 1974).

NECHYBA, T., "Existence of Equilibrium and Stratification in Local and Hierarchical Tiebout Economies with Property Taxes and Voting," *Economic Theory* 10(2) (1997a), 277–304.

——, "Local Property and State Income Taxes: The Role of Interjurisdictional Competition and Collusion," *Journal of Political Economy* 105(2) (1997b), 351–84.

NESHEIM, L., "Equilibrium Sorting of Heterogeneous Consumers Across Locations: Theory and Empirical Implications," Dissertation, University of Chicago, 2001.

OATES, W., "The Effects of Property Taxation and Local Spending on Property Values: An Empirical Study of Tax Capitalization and the Tiebout Hypothesis," *Journal of Political Economy* 77 (1969), 957–71.

O'SULLIVAN, A., T. SEXTON, AND S. SHEFFRIN, *Property Taxes and Tax Revolts* (New York: Cambridge University Press, 1995).

PALMQUIST, R., "Welfare Measurement for Environmental Improvements Using the Hedonic Model: The Case of Nonparametric Marginal Prices," *Journal of Environmental Economics and Management* 15 (1988), 297–312.

——, "Hedonic Methods," in J. Braden and C. Kolstad, eds., *Measuring the Demand for Environmental Quality*. (Amsterdam: North Holland, 1991).

PETRIN, A., "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy* 110(4) (2002), 705–29.

POLINSKY, A., "The Demand for Housing: A Study in Specification and Grouping," *Econometrica* 45(2) (1977), 447–61.

——, AND D. ELLWOOD, "An Empirical Reconciliation of Micro and Grouped Estimators of the Demand for Housing," *Review of Economics and Statistics* 61(2) (1979), 199–205.

POTERBA, J., "Taxation and Housing: Old Questions, New Answers," *American Economic Review* 82(2) (1992), 237–42.

SCOTCHMER, S., "Hedonic Prices and Cost/Benefit Analysis," *Journal of Economic Theory* 37 (1985), 55–75.

——, "The Short-Run and Long-Run Benefits of Environmental Improvements," *Journal of Public Economics* 30 (1986), 61–81.

SIEG, H., V. K. SMITH, S. BANZHAF, AND R. WALSH, "Interjurisdictional Housing Prices in Locational Equilibrium," *Journal of Urban Economics* 50 (2002), 131–53.

SMITH, V. K., AND J. HUANG, "Can Markets Value Air Quality? A Meta-Analysis of Hedonic Property Value Models," *Journal of Political Economy* 103 (1995), 209–27.

TIEBOUT, C., "A Pure Theory of Local Expenditures," *Journal of Political Economy* 64(5) (1956), 416–24.

WALSH, R., "Analyzing Open Space Policies in a Computable Locational Equilibrium Model," Dissertation, Duke University, 2002.

WILLIG, R. D., "Incremented Consumer's Surplus and Hedonic Price Adjustments," *Journal of Economic Theory* 17 (1978), 227–53.