Available online at www.sciencedirect.com
SCIENCE DIRECT
ELSEVIER Journal of Environmental Economics and Management 47 (2004) 559–584
JOURNAL OF ENVIRONMENTAL ECONOMICS AND MANAGEMENT
http://www.elsevier.com/locate/jeem

General equilibrium benefits for environmental improvements: projected ozone reductions under EPA's Prospective Analysis for the Los Angeles air basin

V. Kerry Smith,[a,c,*] Holger Sieg,[b] H. Spencer Banzhaf,[c] and Randall P. Walsh[d]

[a] *Agriculture and Resources Economics, North Carolina State University, Box 8109, Raleigh, NC 127695-8109, USA*
[b] *Carnegie Mellon University and National Bureau of Economic Research, USA*
[c] *Resources for the Future, USA*
[d] *University of Colorado, Boulder, USA*

Received 11 July 2002; revised 27 August 2003; accepted in revised form 28 October 2003

Abstract

This research demonstrates how locational equilibrium models can be used for benefit measurement with the detail required to match EPA's benefit analysis for the first Prospective Analysis. Using the projected changes in ozone concentrations for 2000 and 2010 together with the Sieg et al. (Int. Econ. Rev., forthcoming) estimates for household preferences for housing, education, and air quality, this paper measures general equilibrium willingness to pay for the policy scenarios developed for the Prospective study as they relate to households in the Los Angeles area. Benefits are evaluated taking account (at the household level) of initial air quality conditions, relocation based on changes in ozone, and price changes. The framework generalizes the partial equilibrium/general equilibrium comparisons available with conventional computable general equilibrium and property capitalization models. Estimated general equilibrium gains from the policy range from $33 to $2400 annually at a household level (in 1990 dollars).
© 2003 Elsevier Inc. All rights reserved.

JEL classification: H23; Q28; R13

Keywords: Benefit analysis; Locational equilibrium; Ozone improvement

*Corresponding author. Fax: +919-515-6268.
E-mail address: kerry_smith@ncsu.edu (V. Kerry Smith).

0095-0696/$ - see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/j.jeem.2003.10.006

1. Introduction

Research using revealed preference methods in environmental economics has generally sought to measure the benefits associated with small changes in environmental quality. For example, hedonic property value methods offer measures of the incremental willingness to pay for small changes in site specific amenities. By contrast, the tasks that must be addressed by policy analysts require a framework capable of measuring the benefits from large, often spatially diverse, changes in amenities. In our previous research, summarized in [29], we demonstrated that locational equilibrium models offer a new framework for evaluating these types of policies. In that paper, we developed the basic general equilibrium framework, discussed estimation of the model's structural parameters, and illustrated how it could be used with actual changes in air quality conditions between 1990 and 1995.

In this paper, we expand this line of research by considering EPA's policy alternatives, developed for the evaluation of the 1990 Clean Air Act Amendments and reported in the first Prospective Analysis. The main purpose of this paper is to demonstrate that locational equilibrium models of the type considered in [29] can be used for benefit measurement at the level of detail required for realistic policy assessments. In particular, this paper uses the same projected spatial variation in ozone concentrations as was developed for the Agency's benefit analysis for the LA Air Basin as part of the Prospective study to compute partial and general equilibrium benefit measures.

Our analysis indicates that the estimated annual general equilibrium benefits in 2000 and 2010 associated with the ozone improvements due to continuing the policies mandated under the 1990 Clear Air Act Amendments will be dramatically different by income group and location within the South Coast Air Quality Management District. The gains range from $33 to about $2400 per household (in 1990 dollars). These differences arise from variations in air quality conditions, income, and the effects of general equilibrium price adjustment.

This paper builds on a large literature in environmental economics considering the relationship between partial and general equilibrium welfare measures.[1] Models for comparing partial (PE) and general equilibrium (GE) effects can take a variety of forms. Most of the empirical measures have used static computable general equilibrium (CGE) models that focus on consistent treatment of product and factor markets with boundary conditions that utilize both the zero profit condition for firms and (in static models) zero savings for consumers. For the most part, they have evaluated general equilibrium price effects. Examples of this research have found that large changes in environmental regulations [17], or in climate attributes influencing production [19], can result in appreciable price changes outside the sector(s) directly affected. Hicksian measures of consumer surplus indicated marked differences between the partial and the general equilibrium welfare measures.

[1] Over 30 years ago Harberger [16] observed that "While it is clear that no theoretical obstacle stands in the way of taking such considerations (general equilibrium effects) into account, it is in fact rarely done in studies involving applied welfare economics" (p. 791, parenthetical phrase added). There have been few exceptions to his judgment in applied policy analysis in the 30 plus years since he prepared these remarks. Only the cases cited above, along with efforts to evaluate the welfare costs of introducing standards, permits, or taxes in distorted economies (see [6,14] as examples) and some analyses of trade and the environment [13] have attempted to consider general equilibrium welfare measures.

A second set of models considering PE/GE comparisons has been framed in terms of the capitalization of exogenous policies into land (or property) values.[2] In this case, the analyses focus on: (a) the conditions when rents capitalize the effects of exogenous, location specific, changes in amenities or disamenities; and (b) the relationship between these rent changes and Hicksian welfare measures. To our knowledge there has been no effort, until recently, to use these models to develop numerical comparisons of PE/GE welfare measures.[3]

Most summaries of this second line of research focus on Starrett [32] and Scotchmer [26,27] as providing the most complete analyses of these two questions. However, both make restrictive assumptions that are relaxed in our analysis. For example, to derive the fundamental capitalization relationship for his ''internal capitalization'' model, Starrett assumes the equivalent of weak complementarity and additive separable preferences in the numeraire good. Scotchmer's [27] short run analysis also assumes income effects are negligible, noting in her evaluation of the importance of these assumptions that: ''income effects will be troublesome when income classes are segregated in space…'' (p. 72). Her long run analysis adopts a form comparable to Starrett that is transferable in the numeraire, so heterogeneity in the effects of income on demand for land (or housing) is not important. As a result, her long run corrections to short run benefit measures focus on the changes in population density induced by an exogenous amenity change.

As noted above, this paper follows a different approach, which is based on a class of hierarchical locational equilibrium models first estimated by Epple and Sieg [12]. In Sieg et al. [29], we have shown how to adapt the locational equilibrium framework to estimate partial and general equilibrium benefits of large changes in spatially delineated local public goods and amenities. The approach used in this paper generalizes the PE/GE comparisons available with conventional CGE and property capitalization models in three ways.

First, heterogeneity in household income effects is integral to our analysis of the effects of large-scale policy changes and is used, along with unobserved heterogeneity in household tastes for public goods, to characterize the locational equilibrium. Second, the framework allows for spatial delineation in household adjustments to policy changes at a scale consistent with the requirements of modern environmental policy analyses.[4] Households can adjust to changes in a spatially

[2] These analyses are also sometimes discussed in terms of hedonic property value models or urban spatial structure models. See [22], pp. 108–116 for an early overview of this literature.

[3] Two important papers by Cropper et al., [8] and Cropper et al., [7] use one type of locational sorting model—an assignment framework with an assumed exogenous supply of houses—to evaluate the performance of hedonic and random utility models in measuring marginal willingness to pay for site attributes. They have not considered the effects of large exogenous changes in site amenities for measures of welfare. In unpublished research, Bayer et al. [2] use a different type of locational equilibrium model based on a Berry, Levinsohn and Pakes [4] share inversion framework to consider how policy impacts sorting outcomes and measures of incremental willingness to pay.

[4] The models underlying the Hazilla-Kopp, Kokoski-Smith, and Goulder et al. analyses all omit spatial details. They describe economic activities as if they took place at a single point. This is a common feature of many CGE models. Espinosa and Smith [13] relax the assumption by taking advantage of trade distortions to locate economic activities and to introduce an air diffusion model. It is probably unreasonable to expect significant spatial and sectoral detail in these types of national CGE models. This limits the opportunity to use them for spatially delineated policies. An alternative strategy distinguishing six to eight broad regions with considerable sectoral detail would certainly be consistent with the scale of current international models. Moreover, one could also envision multiple locational equilibrium models for important metropolitan areas. These models could be used to evaluate the potential implications of market (and price) adjustment for benefit measures applied assuming that a partial equilibrium framework was adequate. This approach

delineated environmental public good by moving. As a result, the price an individual pays for housing, his *selected* level of the public good, and the amount of housing demanded are likely to be different from what would have been realized with the baseline community. Thus, the benefit measures reflect endogenous changes in *both* prices and the environmental resources. Third, and finally, there is a consistent accounting of the distributional effects of these large-scale policies, tracking how heterogeneous (in income and tastes for public goods) households gain (or lose) as a result of these policies.[5]

Section two reviews the methodological framework used in this paper. Section three discusses how some of the policy scenarios used in EPA's Prospective Analysis can be evaluated within this framework. Section four summarizes the main findings of our study. In contrast to methods used by the EPA in the Prospective Analysis, our analysis relies exclusively on revealed preference arguments. This feature raises the question of whether the results of this study can be compared to EPA's. We discuss this question and related issues in section five. Section six offers some concluding remarks.

2. The locational equilibrium model

The framework used in this paper for benefit measurement is based on a locational equilibrium model. An important advantage of this framework stems from the model's consistent treatment of preference heterogeneity throughout estimation, computation of equilibrium, and welfare analysis. Our analysis in this paper draws heavily on our earlier research on measuring benefits of air quality improvements in the LA Air Basin. A more detailed description of the methodological issues and a discussion of how the parameters of the model can be estimated with publicly available data sources can be found in [29]. In this section, we provide a brief review of the structure of the model, summarize some of our previous findings as they apply to this study, and describe how the framework can be used to evaluate the policy scenarios developed for EPA's Prospective Analysis.

(*footnote continued*)
parallels the EPA strategy for air quality modeling. That is, separate regional analyses for the Eastern and Western US were undertaken in the Prospective Analysis. The Eastern model had a finer resolution than the Western version. In addition, higher resolution modeling for urban areas was undertaken for Los Angeles, Phoenix, and San Francisco (see [33] Appendix C for details). Thus, one approach to developing insights into general equilibrium effects as a routine part of policy analysis would call for developing locational equilibrium models for each of these metropolitan areas to match the urban scale air quality models and use their results in separate GE assessments for each area. This strategy would allow the evaluation of GE effects to be tailored to local circumstances, including both air quality conditions and housing substitution alternatives.

[5] Of course, to realize these results, there are a number of assumptions. We assume households rent their homes, so incomes do not reflect annualized gains (for homeowners) due to re-evaluation of their initial homes. These rents accrue to absentee landlords. We offer an approximate measure of the annualized size of these "wealth" effects, but incorporation of their impact on the locational equilibrium is beyond the scope of this analysis. Relocation is assumed costless, so we likely overstate the effects of heterogeneity on the differences between partial and general equilibrium welfare measures. Initial research by Hallstrom and Smith [15] indicates it may be possible to partially relax the absentee landlord assumption. However, their results acknowledge that the task of dealing with a market with a mix of renters and owners as well as the nature of the "income" or "wealth" gains (or losses) due to policy effects on existing structures were areas requiring significant additional research.

2.1. Structure of the model

The locational equilibrium model can be described as if households' decisions were undertaken in two stages—the selection of a best community (j^*) and, then, conditional on that choice, an optimal demand for the community specific good.[6] In our application, this community specific good is housing. Eq. (2.1) describes the first stage of this choice process, with m_i denoting the ith household's income and p_j representing the jth community's price for housing. $\theta(q_j, a_j, \varepsilon_j)$ is a separable sub-function in preferences that corresponds to an index, θ_j, of air quality, a_j, and other local public goods, q_j. ε_j is unobserved community specific heterogeneity and α_i is an unobserved taste parameter[7]

$$j_i^* = \arg\max_{j \in A}\{V(\alpha_i, m_i, \theta(q_j, a_j, \varepsilon_j), p_j)\} \quad \forall_i, \tag{2.1}$$

where A = set of communities and $V(\cdot)$ denotes the indirect utility function. By treating α_i as a random variable, jointly distributed with income, the model allows for unobserved heterogeneity in household preferences for public goods. The optimal housing demands for the ith household, h_i, given j_i^* (the optimal community for household i), is given by

$$h_i = -\frac{V_{p_{j^*}}}{V_{m_i}} \quad \forall_i. \tag{2.2}$$

Locational equilibrium requires a set of prices such that market demand for housing equals supply in each community. We omit the subscript i in what follows. Using this continuous formulation, with $f(\alpha, m)$ the joint density for α and m, the market equilibrium is defined by

$$S_j(p_j) = \int_{c_j} h(p_j, m) f(\alpha, m)\, d\alpha\, dm \quad \forall_j, \tag{2.3}$$

where c_j = set of all households such that $j_i^* = j$, $S_j(p_j)$ = supply of housing in community j.

The single crossing property (SCC) is important to characterizing the properties of locational equilibria. It implies that pairs of indirect utility functions, plotted for different income levels in terms of p and θ, will intersect only once. It has been discussed in a number of market contexts associated with the sorting of heterogeneous agents among differentiated goods. In urban location models, Ellickson [10] first discussed the condition as sufficient for perfect stratification by income when this is the only source for preference differences. To our knowledge, SCC has not been used in applications of revealed preference models for non-market valuation. As a result, it may seem to have little relevance to Hicksian welfare measurement. However, this judgment is incorrect. The single crossing property generalizes the Willig [35] condition required to consistently derive Hicksian measures of the value of quality changes from Marshallian demand functions for a

[6] This interpretation is simply for heuristic purposes. The framework implies that optimal choices of housing and the numeraire are evaluated for every alternative and the community yielding the highest utility is selected.

[7] In the empirical implementation of the model the function $\theta(\cdot)$ is assumed to be linear in a measure of local public education and a measure of the ozone concentration. The coefficient of the education measure (q_i) is normalized to unity. Thus, for our empirical application the index of public goods, θ_j, is equal to $q_i + \gamma a_j + \varepsilon_j$, with a_j the measure of the ozone concentration. Sieg et al. [29] also investigate the effects of including other location specific local public goods, such as the crime rate.

private good serving as a weak complement.[8] Our analysis also assumes that the index of public goods can be treated as exogenous so that we can focus on the equilibrium prices in housing markets.

If household preferences are consistent with the single crossing condition then a locational equilibrium must satisfy three conditions: boundary indifference, stratification, and ascending bundles. Boundary indifference refers to the existence of households that are indifferent between two communities that have been ranked to adjacent positions based on their prices and public goods. Stratification implies that households are distributed across communities with an ordering by income (conditional on the taste parameter) and the taste parameter for public goods (conditional on income). Ascending bundles implies both the index of public goods and the index for housing prices increase in the same order.

One of these features, ascending bundles, implies that communities can be ordered by the prices for a standardized unit of housing or by the values of the index, θ, of local public goods. Moreover, these orderings will be the same for all households. As a result, we know that if we consider communities j and $j+1$ ordered by price, these will be the same communities as those ordered to positions j and $j+1$ by θ. Using boundary indifference together with this feature (i.e. setting $V(\alpha, m, \theta_j, p_j) = V(\alpha, m, \theta_{j+1}, p_{j+1})$) we can derive an expression to sort households. That is, substituting the expression for preferences into the definition of boundary indifference and collecting the terms with θ_j, θ_{j+1}, p_j, and p_{j+1} into one expression and the terms describing "unique" households (e.g. α and m) into another, we define a household's threshold for communities in terms of their prices and local public goods.

The Epple–Sieg model recognizes that the sorting to obtain public goods implies features for the observed distribution of income (or housing expenditures) as well as the price indexes. To see this point, consider the indirect utility function (2.4) used in our empirical analysis. This specification is consistent with constant income and price elasticities of demand for housing and separability of community specific public goods from the housing demand, given a community location

$$V(\alpha, m, \theta_j, p_j) = [\alpha \cdot \theta_j^{\rho} + (\exp((m^{1-v} - 1)/(1-v)) \exp((1 - Bp_j^{\eta+1})/(1+\eta)))^{\rho}]^{1/\rho}. \tag{2.4}$$

The boundary indifference condition can be used to isolate those combinations of α and m associated with the points of indifference. Eq. (2.5) provides this expression for the preference specification in (2.4), with the left side of the equation capturing each household's features and the

[8] See [5,30,35] for a discussion of the Willig condition. One form used to express it (for our model) would hold that:

$$\frac{\partial}{\partial m}\left(\frac{dp}{d\theta}\right) = 0.$$

The single crossing condition can also be written in terms of the change in this slope (which is a ratio of partial derivatives of the indirect utility function) with respect to income as:

$$\frac{\partial}{\partial m}\left(\frac{dp}{d\theta}\right) > 0.$$

right side defined from the community attributes

$$\ln(\alpha) - \rho\left(\frac{m^{1-v} - 1}{1 - v}\right) = \ln(\exp(\rho(1 - Bp_{j+1}^{\eta+1})/(1+\eta)) - \exp(\rho(1 - Bp_j^{\eta+1})/(1+\eta))) - \ln(\theta_j^{\rho} - \theta_{j+1}^{\rho}). \tag{2.5}$$

By starting with the lowest ranked community, we can recursively assign households to communities. The left side of Eq. (2.5) provides the index of preferences to use in this sorting. Households whose preferences imply they are below the condition balancing price and the index of public goods given on the right side of the equation will choose to live in community j or below.

Given a distribution to describe all households based on their taste for public goods (α) and income (m) (i.e. $f(\alpha, m)$), a specification of household preferences (e.g. Eq. (2.4)), and housing supply functions in each community, the locational equilibrium maps levels of public goods into housing prices and distributions of households (by income and tastes) across communities. The main insight of Epple and Sieg [12] was to recognize that one can invert this mapping and estimate the parameters of the model using publicly available data. Estimation is based on a two step procedure. First, we estimate housing prices for each community.[9] Given estimated housing prices, ascending bundles implies that we can rank communities by their desirability using housing prices. This ordering establishes a hierarchy among the set of communities in the choice set.

Second, we estimate the structural parameters of the model. If we observed all relevant components of public good provision, estimation would be straightforward. In that case, we could evaluate the boundary indifference conditions given observed levels of public good provision and estimated prices. Hence, we could characterize sorting of households across communities. In practice, this approach is not feasible. We either do not observe all components of public good provision or some of the components are measured with error. Epple and Sieg [12] demonstrated that we can still estimate the parameters of the model under these more realistic assumptions. The key idea is that we can invert the system of community size equations and recover the implied levels of public good provision. This inversion procedure is similar in spirit to the share inversion procedure proposed by Berry [3] for static random utility models or the CCP estimator for dynamic discrete choice models proposed by Hotz and Miller [18].

The inversion of the community size equations yields levels of public good provision that are consistent with the observed community sizes and the structure of the locational equilibrium model given the estimated housing prices. As before, we can then characterize sorting of households across communities and estimate the parameters of the models using an instrumental variable (IV) estimator, which matches selected moments for the quantiles of the observed community specific distributions of housing and income to the ones predicted by the model. We can also match the implied levels of public good provision to the measures observed in the data.

[9] See [28] for comparison of alternative price indexes and Sieg et al. [29] for more details on the estimates used in this application. The price indexes were estimated using the housing sales data weighted so that this transactions sample matches three attributes of the stock of owner occupied housing according to the US Census—the location, the number of bedrooms, and the age of the house. These hedonic models include the number of bedrooms, bathrooms, lot size, building size, age, presence of a swimming pool, presence of at least one fireplace, the average commuting time to work (matched by census block group), and dummy variables for the miles to the coast.

The estimates used for this policy analysis involve seven types of moments: three housing expenditure quantiles (25th, 50th, 75th), three income quantiles (25th, 50th, 75th), and one based on a linear index for the public goods (θ_j). These predictions for the distributions of income and housing expenditures use the recursive relationships based on the boundary indifference condition (Eq. (2.5)), together with the definition of community sizes derived using the appropriate integrals of $f(\alpha, m)$, to identify the model's parameters. Functions of the community ranks are used as instruments to define additional orthogonality conditions. For example, with one instrument for each orthogonality condition and W_j a 7×7 matrix with the instrument used for each of the seven orthogonality conditions along the diagonal for the jth community, zeros elsewhere, and $m_j(b)$, the 7×1 vector of stacked moments for the jth community expressed in terms of the vector of parameters (b) to be estimated, the following equation defines the GMM objective function:

$$\hat{b} = \arg\min_{b \in \mathbf{b}} \left\{ \frac{1}{J} \sum_{j=1}^{J} W_j m_j(b) \right\}^T V^{-1} \left\{ \frac{1}{J} \sum_{j=1}^{J} W_j m_j(b) \right\}, \tag{2.6}$$

V is the covariance matrix for the moments and J is the number of communities.[10]

2.2. The Los Angeles area model

In our previous research we have assembled a unique data set for the LA Air Basin and estimated the parameters of the model. Here we provide a brief summary of some of the key results, since we will use the same parameter estimates in our evaluation of the scenarios developed for EPA's Prospective Analysis. The area selected consists of a portion of five counties west of the San Gabriel Mountains, including parts of Los Angeles, Orange, Riverside, San Bernardino, and Ventura Counties. Census information for 1990 along with data on the sales prices and housing characteristics for virtually all transactions involving private homes in the area between 1989 and 1991 were combined with spatially delineated information on the ozone concentrations, measures of the quality of local public education, and other local characteristics to estimate the model.[11]

Several different price indexes, definitions for the community choice set, and specifications for the set of local public goods were evaluated. Our policy simulations are based on the community definitions offering the greatest spatial resolution for ambient air quality. This formulation takes advantage of a reorganization of the largest school district in our LA metropolitan area. In April 2000, the LA Unified district was decomposed into 11 smaller units. When these districts are combined with the 92 other districts in the area (excluding the original LA Unified district) we have a total of 103 communities for our model.

Sales prices were converted to imputed rents following Poterba [24]. All the estimated parameters for the model were consistent with a priori expectations and statistically significant.

[10] For a more detailed treatment of our specific implementation decisions in estimation as well as parameter estimates derived under differing assumptions, see [29].

[11] Education was measured using standardized test scores for each district using the 1992–93 California Learning Assessment System Grade Level Performance Assessment test. Average math scores were the basis for our measures. They are reported both at the school district and school levels. For the LA Unified sub-districts we used the individual schools to construct the measures for the new school districts.

The three parameters with most direct interest here are the price and income elasticities for housing, estimated at -0.037 and 0.729, respectively, and the estimated effects of ozone on the public good index. This was estimated at -2.923 with asymptotic standard error of 0.674.[12]

The estimated correlation between income and the taste parameter for local public goods was negative and significant. While households with higher incomes have a greater demand for air quality and public schooling, there is also considerable heterogeneity in the housing and income levels observed in each community. The estimated negative correlation reflects this diversity in the distributions of income and housing expenditures.

2.3. Computing general equilibrium solutions for welfare measurement

When large changes in spatially differentiated public goods are modeled, some assumption must be made about whether household adjustment is allowed. As noted at the outset, the specific results derived in [27] for capitalization and long run benefit measures assume quasi-linear preferences in the numeraire good with equal, invariant, income parameters across households. As a result, in Scotchmer's example, "…the dispersion of population depends only on air qualities (her locational public good) and not on the distribution or amount of private good" (p. 76, clause in parentheses added). Starrett's assumed preferences with his internal capitalization model take a similar form. The Epple–Sieg framework allows for heterogeneity in tastes for public goods and allows income to differentially influence the community specific private good consumed.

Market prices for housing in each of the communities in Epple–Sieg reflect the effects of adjustment. The model is closed with a housing supply function. In our application it is calibrated to the initial equilibrium. We will assume an inelastic housing supply for the policy analysis discussed in the next two sections, but this is not essential to applications of the model (see [15,34] for alternative treatments of supply with the model). It is not possible to solve the model analytically. Therefore, to develop an analysis of the general equilibrium effects of an exogenous change in community specific public goods, we simulate the re-sorting implied by the model's equilibrium conditions (i.e. as outlined in Eqs. (2.1)–(2.3)).

The simulation generates draws of the income and the taste parameter to match their estimated joint distribution and allocates these vectors to each community based on the ordering implied by the boundary indifference conditions (Eq. (2.5)) with the estimated preference parameters. This first step takes estimates for the preference parameters as given, including the joint distribution of income and the taste parameter for public goods. One million pairs of α and m are drawn consistent with the maintained assumption of a joint log normal distribution. We then compute the levels of public good provision in the alternative scenarios that we evaluate. These levels reflect the changes in air quality associated with each of the scenarios from EPA's Prospective Analysis and are consistent with our estimates of unobserved community specific amenities (as discussed in

[12] The absolute magnitude of this parameter offers an estimate of the implied marginal rate of substitution between education and ozone concentrations. It is difficult to interpret due to the differences in the measures used for education (e.g. an average test score) and ozone (an average of the ambient concentrations for the 30 highest readings in a year in parts per million). For this reason, it would seem a more direct basis for evaluating the estimated magnitude of the effect of air pollution would be based on the benefit measures derived using the model.

[29]). A new equilibrium is obtained when we have found a vector of housing prices, which implies housing market clearing in all communities. This simulation offers an opportunity to track gains and losses, movement, and partial versus general equilibrium effects by keeping track of initial and final community assignments for each "simulated household."

Given the hierarchical structure of the model, we can reduce the problem of computing equilibrium to a one-dimensional search. The basic idea is to pick a price in the lowest community and then constrain the prices in all other communities such that housing markets in the J-1 communities clear. Thus, we only iterate on the price of the first community until we find market clearing in the last community.

This process allows the definition of partial and general equilibrium measures of the willingness to pay for the policy giving rise to the exogenous change in ozone throughout the Los Angeles basin. Eq. (2.7) defines, implicitly, the general equilibrium willingness to pay (WTP_{GE}) and (2.8) the partial equilibrium measure (WTP_{PE}) for an individual household

$$V(\alpha, m - WTP_{\mathrm{GE}}, \theta_k^*, p_k^*) = V(\alpha, m, \theta_j, p_j), \tag{2.7}$$

$$V(\alpha, m - WTP_{\mathrm{PE}}, \theta_j^*, p_j) = V(\alpha, m, \theta_j, p_j). \tag{2.8}$$

The symbols used to identify the public good and price in Eq. (2.7) change from (θ, p) to (θ^*, p^*). The subscripts also change from j to k. The asterisk from θ to θ^* designates the exogenous change in air pollution (ozone concentrations in our case). The change from p to p^* identifies that, with household adjustment, the community specific housing prices will change. Finally, the subscript differences provide a reminder that this change is the result of household adjustment. So, the change in air quality at location j is not what is relevant to the household that was located in community j unless the household remains in that community. Rather, it is what the household selects given the new choice set. Thus, the distinction between GE and PE welfare measures is direct. The first allows for price change and incorporates household relocation. The second defines a counterfactual that would not be observable in a world with costless adjustment. It assumes the household experiences the change from θ_j to θ_j^* and prices remain unchanged.

Our analysis assumes other local public goods do not change as the income distributions in each community change with re-sorting. Epple et al. [11] have demonstrated it is possible to incorporate these types of endogenous changes in public goods using a median voter framework, along with a specified production relationship, and a community budget balancing condition through local taxation. However, this type of extension is beyond the scope of this research.

3. Using locational equilibrium models for policy scale benefit assessment

A principal objective of this paper is to demonstrate that locational equilibrium models can be developed at a scale that accommodates the detail required by benefit analyses for actual environmental policy alternatives. To meet this goal we used the scenarios developed for EPA's first Prospective Analysis [33] to provide the changes in ambient ozone concentrations that are evaluated with the model developed in [29].

In the remainder of this section we describe how EPA's detailed simulation of ozone conditions in the Los Angeles area under the different regulatory scenarios were adapted to conform to our model of households' locational choices in this same area.[13]

3.1. Background on the Prospective Analysis

The scenarios for EPA's Prospective report consider the regulatory restrictions assumed to characterize situations without and with Clean Air Act Amendments. These specifications lead to estimates of ambient air quality conditions under these two potential regulatory regimes for 2000 and 2010.

The EPA analysis identified five major source categories for air pollutants: industrial point sources, utilities, non-road engines/vehicles, motor vehicles, and area sources.[14] For each source both a base year level of emissions and projected growth of the specific pollution generating activities were developed for 2000 and 2010 in the absence of CAAA requirements. These projections were modified to reflect control assumptions in each of these two "future years." For volatile organic compounds (VOC) and nitrogen oxide (NO_X), both contributors to ambient ozone, the national assessment estimates a 27% reduction due to CAAA's effects on VOCs in 2000 and a 35% reduction in 2010. For NO_X, the reduction is comparable for 2000 (i.e. 26%) and a little larger in 2010 (39%).

Sectoral emission projections were disaggregated to the state/industry level or below.[15] Spatially differentiated emission estimates, at least at the county level, were then introduced into one of several air quality models. Ozone concentrations were estimated using the Urban Airshed Model (UAM) and the variable grid UAM (UAM-V). For the Los Angeles area, the EPA analysis supplemented the regional scale modeling results with higher resolution analysis. As a result, ambient concentrations on an hourly basis are available at a 4-km resolution for this area.

The EPA air quality modeling for the Los Angeles area also relied on input data from the South Coast Air Quality Management District (SCAQMD).[16] The EPA projections with and without CAAA considered the conditions for two separate 3-day periods (June 23–5, 1987 and August 26–28, 1987) with baseline conditions augmented by the emission profiles associated with each scenario (see Table C-4 in [33] for a summary of emission totals for each scenario in Los Angeles).

[13] Section 812 of the 1990 Clean Air Act Amendments (CAAA) requires periodic assessments of the benefits and costs of the regulations intended to reduce the concentrations of criteria air pollutants. The first Prospective Assessment [33] estimated the annual human health benefits (in 1990 dollars) of Titles I through V of the CAAA to be $68 billion in 2000 and $108 billion in 2010.

[14] EPA's definition for the source categories are given as follows:

(a) Industrial point sources—boilers, cement kilns, process heaters, turbines.
(b) Utilities—electricity producing utilities.
(c) Non-road engines/vehicles—air craft, construction equipment, lawn and garden equipment, locomotives, and marine engines.
(d) Motor vehicles—buses, cars, trucks.
(e) Area sources—agricultural tiling, dry cleaners, open burning, wildfires

[15] Appendix A of EPA [33] provides a detailed description of the models used in the sectoral emission analysis.

[16] The SCAQMD was also the source for the monitoring data used to estimate our economic model, discussed below in section three.

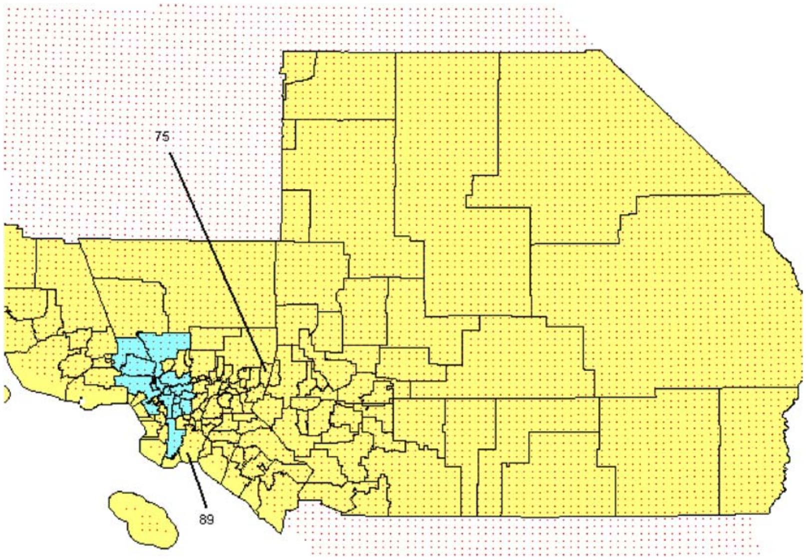

Fig. 1. LAAB school districts and EPA UAM locations for estimated ozone readings. Long Beach Unified is ID #89; Claremont Unified is ID #75.

With the assistance of EPA staff and its contractors, the three analyses undertaken for Southern California were interpolated to a consistent 5-km $\times$ 5-km grid cell pattern for the study area included in our analysis. For each scenario, the latitude and longitude of the centroid for all cells in Southern California were computed. Hourly ozone values were summarized for each grid cell with 10 hourly ozone values (measured in parts per million) as the 5th percentile through the 95th percentile of all the modeled hours in the 6 day simulation period.

Our analysis had access to the results of three sets of simulations from the air diffusion models: (1) baseline runs for ambient ozone concentrations in 1990; (2) ambient concentrations in 2000 and 2010 when air pollution regulations are "frozen" at federal, state, and local controls corresponding to their 1990 levels of stringency and effectiveness; and (3) concentrations in 2000 and 2010 when federal, state, and local rules promulgated under the 1990 CAAA are implemented.

Fig. 1 indicates the geographic scale of these simulations in relation to the school districts, which provide the lowest level of spatial resolution for our model. The "dots" falling within the boundaries of each school district identify the number of projected ozone distributions from EPA's air diffusion analysis. Each is the centroid of a 5-km grid. Our estimated model relied on the average of the highest thirty hourly ozone readings for a year, which does not have an exact counterpart in the available data. As a result, we used the average across grids in a school district for the 0.95 deciles. To assure consistency between the monitored and simulated readings, we constructed school district specific scale adjustments based on the average monitored ozone readings in 1990 for each school district relative to the average of EPA projected concentrations for the 0.95 decile in 1990.[17]

[17] This process parallels EPA's practice in using their simulations (see [33], Appendix C, p. C-25).

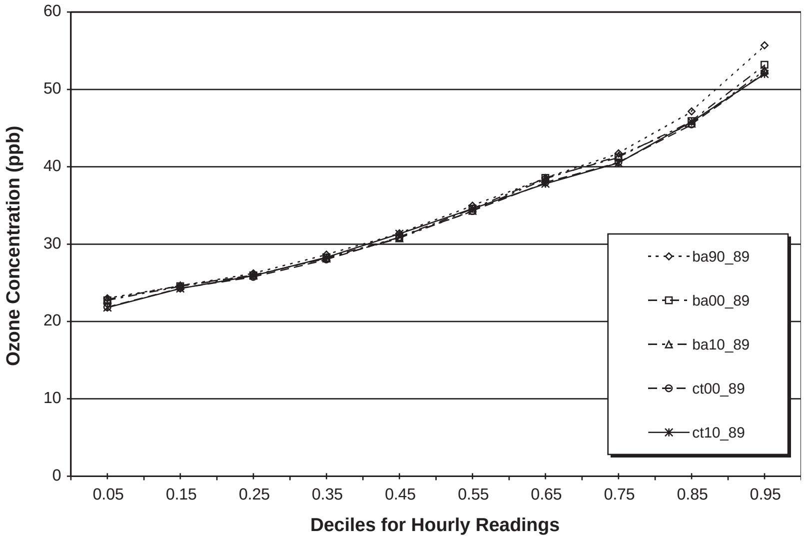

Fig. 2. A comparison of projected ozone concentrations with and without CAAA for Long Beach Unified school district (89).

3.2. Two examples of changes in ozone conditions

Fig. 1 labels two of the school districts in our revised choice set with identifying numbers. Figs. 2 and 3 present the average simulated concentrations (before the adjustment described above, in parts per billion of ozone).[18] Each graph provides five empirical distribution functions, distinguished by year (90 for 1990, 00 for 2000, and 10 for 2010) and air pollution control scenario (baseline without, ba and with CAAA controls, ct). These two cases illustrate a more general pattern. There is a substantial diversity in conditions across school districts. In some cases these empirical distributions do not shift with either the year or control conditions. In others, there are pronounced differences between ba and ct, but they vary by year. For example, considering the 0.95 decile (the right-most point on each graph) in Fig. 2 for the Long Beach Unified school district (ID = 89) (see the lower left side of Fig. 1 for its location), all the computed values for the ozone concentration are relatively constant and below the current 0.12 primary standard (parts per million). By contrast, if we consider a district away from the coast, such as Claremont Unified (ID = 75) in Fig. 3, the picture is quite different with large changes in the ozone concentrations across scenarios.

4. Results

To consider the implications of changes in ozone conditions for housing prices we plot the computed housing price index for a standardized unit of housing and measured index for public

[18] We re-scaled by 1000 to make the distinctions easier to display. Ordinarily ozone is measured in parts per million. These figures are for the readings prior to the re-scaling described earlier.

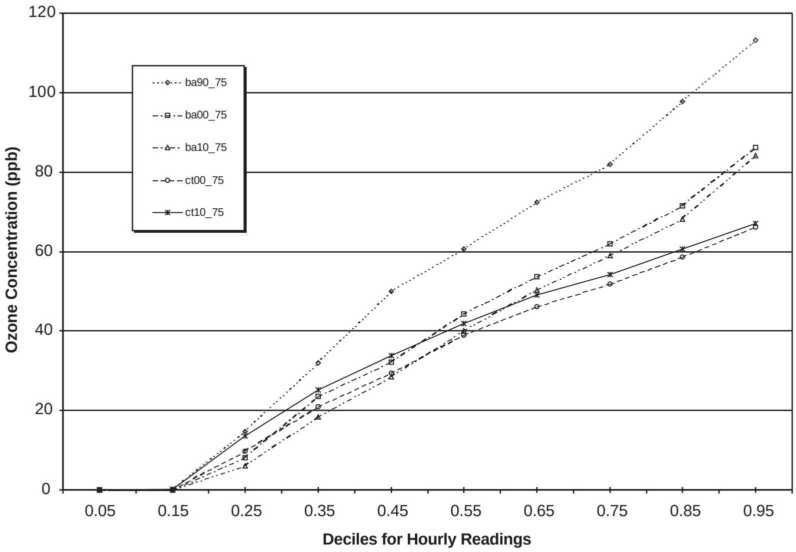

Fig. 3. A comparison of projected ozone concentrations with and without CAAA for Claremont Unified school district (75).

goods for the GE solutions for baseline 2000 and control 2000 cases in Fig. 4. The simulated improvement in ozone conditions in 2000, comparing the baseline and control solutions, lead to substantial price change and household adjustment.[19] The vertical axis plots the housing price index and the horizontal axis has the $\theta = \theta(q, a, \varepsilon)$ index of public goods. Recall the price index used to estimate the model was derived from the community specific fixed effects in hedonic price models that included variables to take account of each home's structural characteristics, proximity to the coast, and commuting time. Thus, these computed GE prices based on the estimated model measure how this type of price index would change with changes in ozone conditions. The θ index combines the education scores and ozone concentrations in each school district using the parameter estimates from the model as weights. Each point in the graph corresponds to a school district. The first line labeled "ba" corresponds to the general equilibrium prices and index of public goods derived after households re-sort from their initial 1990 community location to reflect the ozone conditions EPA's Prospective Analysis estimates for 2000 with no new regulations beyond those in place in 1990. The "ct" line uses the ozone conditions simulated for the control case in 2000.

[19] There are two steps in developing each set of estimates. First, beginning with the 1990 equilibrium allocation of households used to estimate the model, we solve for equilibrium prices and sort households for the baseline 2000 based on the projected ozone distribution, adjusted using the 1990 scaling factor. To compute the 2000 control we assign the control ozone conditions for 2000 to each school district, adjusted using the scale factor (i.e. the ratio of actual 1990 to projected 1990 measures described earlier), and re-solve the model for the new prices and household allocations among communities with CAAA controls and the revised ozone concentrations. This process is repeated to derive the 2010 baseline and 2010 control solutions. The GE solutions for prices at given values of the public good index for baseline 2010 and control 2010 display a similar pattern and are not reported here.

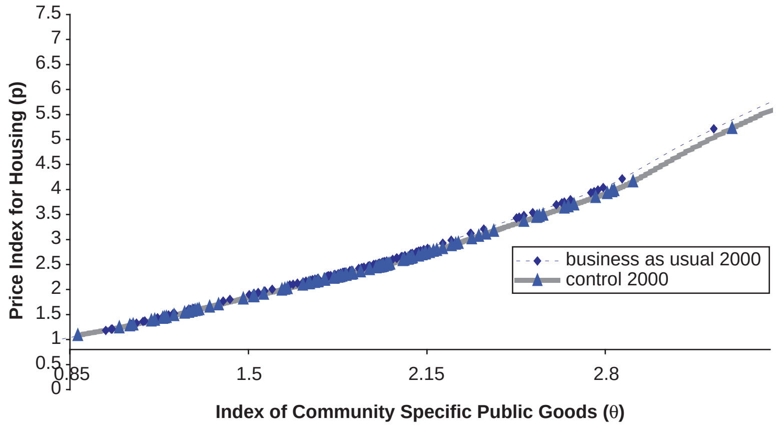

Fig. 4. Computed general equilibrium prices for business as usual and control 2000 (the legend follows the format described in the text).

Beginning with the baseline 2000 household distribution, we compute the re-allocation of households and the implied general equilibrium prices. Only ozone concentration is changing to induce this pattern. Our measures for the performance of local public education are assumed to be constant at the 1990 levels in each school district. Nonetheless, the ordering of the school districts by price and public goods can change with the diverse set of changes in ozone conditions and induced price adjustments. As a result, horizontal or vertical comparisons of these two lines are not informative of changes in the conditions for a specific community. Fig. 4 does clearly indicate that controls lead to a superior choice set over baseline conditions. That is, taken as a whole for nearly all but the lowest ranked communities, it is possible to get improved air quality conditions at the same or lower prices.

4.1. General and partial equilibrium benefit measures

Table 1 illustrates some of the benefit calculations that are possible with the model. The complete WTP distribution is estimated based on all the households that are assumed part of the LA area housing market. Each benefit measure, both the GE and PE, varies across households. To develop summary statistics we must average in some way. We have computed these averages for the households (i.e. pairs of α and m) initially assigned to each school district. More formally, the average WTP^{j}_{GE} for school district j is defined by

$$WTP^{j}_{\mathrm{GE}} = \int_{c_j^0} WTP_{\mathrm{GE}}(m,\alpha) f(\alpha, m)\, d\alpha\, dm / p(c_j^0), \tag{4.1}$$

where c_j^0 = baseline community designation, $p(c_j^0)$ = measure of number of households in c_j^0

Table 1 also reports the averages for: the area as a whole (in the first row under each scenario), each county, and then for two of the school districts we profiled in describing the projected changes in ozone concentrations. Our policy comparison solves the locational equilibrium model under the 2000 and 2010 business-as-usual ozone levels (ba00 and ba10). The 2000 and 2010

Table 1
Alternative benefit estimates with (ct) and without (ba) CAAA regulations for 2000 and 2010

	$\Delta ozone$[a]	Actual[a]		Initial[a]					
		Δp	$\Delta\theta$	$\Delta\tilde{p}$	$\Delta\tilde{\theta}$	WTP_{GE}	WTP_{GE}/m	WTP_{PE}	$\Delta Rent$
2000									
Area-wide	−0.124	0.001	0.027	0.002	0.027	798	0.012	708	−87
County									
Los Angeles	−0.114	−0.005	0.022	−0.004	0.022	933	0.013	819	−111
Orange	−0.134	−0.007	0.019	−0.011	0.016	869	0.013	534	−335
Riverside	−0.154	0.042	0.059	0.046	0.062	179	0.006	517	342
San Bernardino	−0.179	0.045	0.066	0.045	0.066	238	0.007	613	379
Ventura	−0.097	−0.009	0.019	−0.008	0.020	497	0.011	344	−154
Selected school districts									
Long Beach (89)	−0.046	−0.022	0.006	−0.026	0.003	499	0.011	49	−452
Claremont Unified (75)	−0.019	−0.011	0.017	0.028	0.053	607	0.010	1151	595
2010									
Area-wide	−0.105	0.000	0.025	0.000	0.025	789	0.012	723	−59
County									
Los Angeles	−0.094	−0.006	0.020	−0.002	0.023	941	0.012	940	5
Orange	−0.106	0.010	0.015	−0.016	0.011	821	0.012	413	−402
Riverside	−0.183	0.063	0.080	0.059	0.076	186	0.006	547	369
San Bernardino	−0.143	0.028	0.053	0.029	0.054	270	0.008	487	219
Ventura	−0.094	−0.012	0.016	−0.026	0.004	539	0.011	21	−505
Selected school districts									
Long Beach (89)	0.061	−0.031	−0.003	−0.041	−0.012	488	0.011	−236	−720
Claremont Unified (75)	−0.142	−0.031	−0.003	0.005	0.030	481	0.011	524	76

[a] Proportionate changes.

equilibria are computed starting from the ba cases under the assumption that the 1990 CAAA regulations are implemented in the years 2000 and 2010 (ct00 and ct10).

The first column reports the average proportionate reduction in ozone implied by replacing the baseline 2000 (and 2010) "without CAAA" with the projected "with CAAA controls" for the initially assigned school district ($\Delta ozone$). Δp and $\Delta\theta$ in columns 2 and 3 (and labeled "actual") correspond to the proportionate changes in prices and public goods experienced by households whose initial assignment was the identified county or school district. Thus, they reflect the result of any movement between school districts (e.g. $(p_k^* - p_j)/p_j$ and $(\theta_k^* - \theta_j)/\theta_j$, with j the initial community and k the post ozone change community allowing for equilibrium price adjustment). $\Delta\tilde{p}$ and $\Delta\tilde{\theta}$ in the next two columns measure what happened to prices and the public good index in the original location, measured as proportionate changes. The remaining statistics correspond to the average value for general equilibrium willingness to pay (as defined in Eq. (2.7)), the general equilibrium willingness to pay relative to income, the partial equilibrium willingness to pay (Eq. (2.8)), and a measure of the changes in rent, $\Delta Rent$ (i.e. $(p_j^* - p_j) \cdot h_j$).

Averaged over all school districts, the discrepancy between the GE and PE measures of willingness to pay is small, under \$100. This result is not surprising because the analysis holds the number of households constant. They can re-sort among a finite set of alternatives under the changed conditions. Improvements in the amenity conditions at one location lead to movements to that area (adjustment is assumed costless), prices adapt, and the balancing effect of the market reduces the gains available.

Two aspects of these results are especially interesting. First, as we average at lower levels of aggregation, the heterogeneity in gains and losses becomes more apparent. Differences between general and partial equilibrium measures can also increase. Consider, for example, the average for households initially in Orange County for the 2000 scenario. Ozone is projected to decline by 13% on average. This improvement is estimated to yield a \$534 annual gain (in 1990 dollars) considering only the ozone reduction. But, household adjustments lead to housing price declines as well which imply the GE benefits are larger. Of course, the reverse is also possible. In the Claremont Unified school district there is a \$500 difference between partial and general equilibrium measures of ozone improvement. The air quality change is small, about 2%, but housing prices increase by nearly 3%. Thus, the composite effects of both lead to smaller GE benefits.

It is worth noting that in 2010 a small increase in ozone concentration estimated for Long Beach does not necessarily spell a GE loss for the households, because the same policy improved conditions elsewhere. The households initially assigned to Long Beach can move to better ozone conditions and may improve the prices they pay for housing. This composite explains the difference for this school district in 2010.

The size of our estimates for the rent changes might seem to offer a simple explanation for the difference between the partial and general equilibrium welfare measures. That is, adding the average rent change to the PE measure in most cases leads to a close approximation to the GE welfare measures. However, this outcome is not one that can be expected to hold generally. To see this point we can rewrite Eqs. (2.7) and (2.8) in terms of expenditure functions as

$$WTP_{\mathrm{PE}} = e(p_j, \theta_j, V^0, \alpha) - e(p_j, \theta_j^*, V^0, \alpha), \tag{4.2}$$

$$WTP_{\mathrm{GE}} = e(p_j, \theta_j, V^0, \alpha) - e(p_k^*, \theta_k^*, V^0, \alpha). \tag{4.3}$$

The difference, $d = (WTP_{\mathrm{GE}} - WTP_{\mathrm{PE}})$, is Eq. (4.4). This compares the income required to realize the utility designated by V^0 in the old community with new air quality and old prices versus what is required in the new community with new air quality and new prices

$$d = e(p_j, \theta_j^*, V^0, \alpha) - e(p_k^*, \theta_k^*, V^0, \alpha). \tag{4.4}$$

There is no reason to expect d will equal $\Delta Rents = (p_j^* - p_j) \cdot h_j$, with h_j the Marshallian demand at (p_j, θ_j) as defined in general terms in Eq. (2.2). Our rent measure, $\Delta Rents$, is somewhat similar to the rent change Lind [20] used to define an aggregate upper bound on benefits from an exogenous change in a spatially delineated public good. This parallel arises because our measure focuses on the initial housing in each community (h_j) (similar to his focus on parcels rather than households). It measures how h_j is re-valued after the price change (i.e. $(p_j^* - p_j)$). It does not take account of the full adjustments in the amount of housing purchased by each household in response to new prices and public goods that are implied in the second term in d. Because $e(\cdot)$ is dominated by housing expenditures, small changes in θ (from θ_j to θ_j^*) will allow $p_j h_j$ to approximate $e(p_j, \theta_j^*, V^0, \alpha)$. Thus, in the absence of a clear expectation for $e(p_k^*, \theta_k^*, V^0, \alpha)$ in relation to $p_j^* h_j$ it is not possible to conclude that $WTP_{\mathrm{GE}} \approx WTP_{\mathrm{PE}} + \Delta Rent$.

4.2. Extent of household adjustment

Our simulation strategy identifies each generated (α, m) pair as a "household," assigns them to an initial community to determine the baseline housing prices, and then sorts these pairs in response to the specified exogenous changes in ozone concentrations until the new equilibrium prices provide no incentive for relocation. One gauge of the importance of the assumption of zero adjustment costs, suggested by a referee, is the fraction of households moving in response to the policy change being evaluated. Table 2 provides a summary of the percentage of households moving for five illustrative school districts in response to the air quality regulations impact on ozone in 2000. We selected the lowest and highest income school districts to highlight a point discussed in the next subsection. At the boundaries of the set of communities defined in terms of price, there is less scope for adjustment. The remaining three communities were selected to illustrate cases where everyone leaves their initial location (LA Unified B), the majority stay (Long Beach), and the majority, but not all, leave (Glendale Unified).

The table reports the percent of the initial household selecting each of the identified school districts as their new location. The numbers in bold are the own community selections. Dashes indicate percentages below one percent. Rounding error accounts for the percentages in each group diverging slightly from 100%.

Three points should be noted in reviewing the results for these illustrative communities. First, if households are not mobile, prices will not adjust. Second, when all the simulated households

Table 2
Household relocation in response to CAAA regulations for 2000[a]

School district	Percentage relocating to other school districts[b]																
	31	89					113					84				69	
	31	58	70	86	89	100	14	70	72	75	81	9	84	103	110	69	101
Baseline location																	
San Jacinto Unified (31)	**100**																
Long Beach (89)		5	5	15	**67**	7											
LA Unified (113)							41	—	40	—	18						
Glendale Unified (84)												31	**43**	10	17		
Beverly Hills (69)																**100**	—

[a] The school districts receiving households from the five baseline locations (in the ba 2000 scenario) in response to the price and ozone changes associated with the 2000 control (CE) scenario are: 9 Los Alamitos Unified, 14 Santa Ana Unified, 58 Ojai Unified, 70 Bonita Unified, 72 Centinela Valley Union High, 75 Claremont Unified, 81 El Monte Union High, 86 Hacienda La Puente Unified, 100 Rowland Unified, 110 Las Virgenes Unified.

[b] The first line identifies the identification number for the originating school district. The second line identifies the new locations. A dash implies values less than one percent. The numbers in bold identify the percent of households remaining in their original community.

initially assigned to a community leave, this outcome does not mean the school district is empty. The model requires prices to adjust with demand for housing equal to the initial supply (i.e. Eq. (2.3) is satisfied in equilibrium). Thus, the baseline supply is occupied in the new equilibrium with new simulated households.

Finally, in the absence of information on the extent of migration, we believe it is difficult to interpret these measures.[20] They do have a value in helping to understand how both price and public good conditions influence household adjustment.

4.3. *Distributional effects of policy*

Table 3 summarizes the type of distributional information that can be developed using the model. For the two simulation years we present the general equilibrium willingness to pay associated with the 25, 50, and 75 percentile in each of six school districts as well as the share of income that this represents for each quartile. We ranked school districts by average income from lowest (ranked 1) to highest (ranked 103). To illustrate the possibilities for assembling distributional information, the gains are presented for the first community, the 20th, 50th, 70th, 90th, and highest average income communities. The WTP distributions relate to the distributions for each community and reflect the heterogeneity in income and tastes for public goods (α) within each location. Notice that the lowest and the top ranked communities do not change. Those in the middle do. We also find that the households in the lowest ranked community lose despite improved air quality. There is large price appreciation in both 2000 and 2010.

Gains are a modest fraction of income but are diverse across the households sorted into these communities as we would expect. The highest income community gains the most in absolute terms because it experiences a reduction in ozone and a reduction in housing prices.[21] This outcome is due to the fact that Beverly Hills has the highest prices in baseline conditions and, thus, an improvement in substitute communities' situation will reduce its housing prices.

5. Discussion

Our findings suggest that the locational equilibrium model has the potential to be used as part of a benefit analysis of environmental policy alternatives. This claim raises the question of whether the benefit estimates presented in this paper should be compared to EPA's measures. The methodology used in most EPA benefit-cost analyses and in developing the benefit measures reported in the Prospective Analysis is completely different from what we used. EPA's analysis

[20] Some reviewers of the Epple–Sieg [12] framework have suggested the Bayer et al. [2] model avoids the questions posed by 100% relocation and as a result offers a superior modeling strategy. It is important to recognize that the reason this framework does not display 100% relocation stems from their definition of market equilibrium. Bayer et al. define equilibrium prices for homes by a probabilistic requirement for the vector of housing prices. They must assure the sum of the probabilities that each house will be occupied is unity. These conditions imply none of the simulated households in their policy analyses actually are located in a house or a neighborhood. Each household has a vector of probabilities for feasible houses. As a result, relocation cannot be defined in this framework. Policies simply change the vector of probabilities for each household.

[21] This analysis assumes absentee landlords. For homeowners, these gains would be offset by the drop in rents.

Table 3
Willingness to pay distributions: effects of CAAA policy for 2000 and 2010

School district	Average income	Actual[a]			Distribution of GE gains[b]				
		Δp	$\Delta\theta$	WTP^{25}_{GE}	sh(25)	WTP^{50}_{GE}	sh(50)	WTP^{75}_{GE}	sh(75)
2000									
1. San Jacinto Unified (31)	19,602	0.068	0.068	−59	−0.003	−40	−0.002	−28	−0.002
20. Corona-Norco (20)	37,001	0.024	0.051	205	0.008	279	0.009	373	0.010
50. LA Unified B (113)	51,918	0.010	0.036	357	0.010	484	0.011	644	0.012
70. LA Unified A (112)	67,129	0.009	0.017	506	0.011	689	0.012	921	0.013
90. Glendale Unified (84)	100,174	0.007	0.031	940	0.013	1275	0.015	1694	0.017
103. Beverly Hills (69)	353,200	−0.003	0.017	3899	0.016	5331	0.018	7406	0.020
2010									
1. San Jacinto Unified (31)	20,369	0.125	0.112	−144	−0.008	−105	−0.006	−74	−0.005
20. Corona-Norco (20)	36,430	0.020	0.049	225	0.009	306	0.010	412	0.011
50. Ojai Unified (58)	50,828	−0.011	0.016	337	0.009	457	0.010	605	0.012
70. Placentia-Yorba Linda (12)	68,298	0.001	0.024	507	0.011	678	0.012	894	0.013
90. Las Virgenes (110)	101,736	0.013	0.034	924	0.013	1262	0.014	1676	0.016
103. Beverly Hills (69)	346,928	−0.001	0.024	4611	0.019	6370	0.021	8769	0.024

[a] Δp and $\Delta\theta$ are proportionate changes.
[b] sh(j) designates the quantile as a fraction of household income.

relies on a damage function approach. That is, the Prospective Analysis estimates the health effects of air pollution changes using concentration/response functions (i.e. the damage functions) from the epidemiological literature, and values the estimated effects by transferring per-unit values from contingent valuation studies, or, for mortality, average values of a statistical life (VSL) measured based on labor market studies. In contrast, our approach begins with the preference specification used to develop the estimates in [29]. Changes in air pollution are assumed to influence location specific public goods. As a result, the benefit measure defined for reductions in air pollution follows directly from our earlier estimates of this preference function (e.g. as given in Eqs. (2.7) and (2.8)).

Unfortunately, the EPA report does not provide welfare estimates disaggregated by region. The report also considers several criteria air pollutants (particulate matter (PM10 and PM2.5), ozone, carbon monoxide, nitrogen oxides, and sulfur oxides). In the EPA analysis, the simple US per-household national average benefit for the 2000 ba scenario was \$578, compared to \$711 from our study for Los Angeles.[22] If the household average, computed from their national analysis, was

[22] This is the area-wide WTP_{GE} from Table 2, net of rent transfers. Using the aggregate estimates for all pollutants and regions for regulations associated with Titles I through V (e.g. those related to the criteria pollutants), EPA's estimated annual benefits for 2000 (in 1990 dollars) were \$71 billion. This estimate corresponds to EPA's central estimate and includes avoided mortality (63 billion), avoided morbidity (5.1 billion), and ecological and welfare effects (3 billion). In 2000, the US population was reported at 282.4 million people. Assuming 2.3 persons per household, this would yield 122.8 million households and an annual benefit per household of \$578 for the emission reductions in 2000 implied by CAAA regulations compared to holding emission restrictions at their 1990 levels but allowing economic growth to 2000.

relevant for the Los Angeles area, it might be argued that EPA's estimate should be larger due to the larger number of air pollutants involved. In the absence of an analytical basis relating the damage function and locational equilibrium/revealed preference methods, it is not possible to assess this conjecture or to provide specific reasons for differences in these estimates. Nonetheless, the relative merits of each approach to policy analysis deserve discussion.

There is also precedent for such comparisons. Most of EPA's early benefit–cost assessments for regulations compared estimates derived using different methods. For example, the assessment of the national ambient standard for particulate matter prepared in the early eighties, reported damage function estimates distinguishing mortality, acute and chronic morbidity, household soiling, and materials damage as separate categories assumed to affect households. Benefit measures for the same proposals using hedonic property and wage results were also reported (see Table 1–10, [21]).[23] Different aggregation schemes were considered to attempt to gauge the difficulties due to double counting that might be posed by these comparisons. The potential for duplication was acknowledged to be a problem for both the damage function estimates and the revealed preference measures. Overlaps in the damage function were discussed in terms of mortality and morbidity effects.

With that background, consider first our approach. It has the advantage of taking account of general equilibrium effects. It also explicitly accounts for the distribution of welfare effects, recognizing both household heterogeneity in taste for public goods and the effects created by general equilibrium price changes.

By contrast, one might argue that the damage function approach is easier to apply. This advantage applies to benefit transfers generally, not the damage function approach in particular. As pointed out by Desvousges et al. [9], it is as feasible to transfer estimates of revealed preference values for air quality per se, without computing intermediate physical damages. Indeed, this paper illustrates how this task might be accomplished by transferring a structural model, capable of developing general equilibrium benefit estimates, *outside* the sample context used for its estimates.

We have transferred our results only to Los Angeles (the same area where the model was estimated). With the assumption of similar distributions of tastes across metropolitan areas, and correlation between tastes and incomes, it might be possible to extend our transfer to different metropolitan areas. This proposal requires many assumptions, but such conditions are also true for transfers of epidemiological and other approaches to computing the willingness-to-pay for health improvements (i.e. mortality and morbidity effects).

There are other important differences between the two methods. The damage function approach makes limited assumptions about households' behavioral responses to pollution. In contrast, our model assumes households can make costless, locational adjustment. It seems reasonable to expect that our assumption overstates the importance of general equilibrium effects. With household adjustment costs, responses would be smaller. While these costs probably cannot be incorporated into the equilibrium model, they could easily be incorporated into the welfare simulations. Any level of adjustment costs would lie between the no adjustment-cost case (our general equilibrium measure) and the infinite adjustment-cost case (our partial equilibrium measure).

[23] For a recently published discussion of this comparison along with other policy and litigation uses of hedonic models, see [23].

Another common critique of revealed preference methods based on locational decisions is skepticism that households really sort by air quality differences. In fact, at least in fairly polluted cities such as Los Angeles, such sorting is indicated qualitatively with another market test, the market for information. Newspapers and other media report detailed air quality information daily; real estate agents report frequent interest in neighborhood air quality from new buyers. Conceding this point in principle, a weaker version of the critique may claim only that the sorting responses due to air quality differences are likely to be so small that they are effectively lost in the statistical noise of many spatial differences. We acknowledge the challenge posed by estimation, but caution against ignoring similar difficulties in the damage-cost approach. Epidemiology also faces such difficulties, with important unobservables, such as indoor pollution, correlated both spatially and temporally.[24]

One specific statistical issue is the problem of collinearity among pollutants. As with other revealed preference work, our model addresses this problem by using only one air pollutant. This decision raises the question of interpreting the one pollutant as a proxy for all correlated air pollutants. Future work should address this question, and consider the use of explicit aggregates such as the EPA's Air Quality Index. By contrast, one of the advantages of the damage-cost approach is its ability to handle multiple pollutants, including minor pollutants unlikely to cause sorting. However, the separate effects attributed to each pollutant also rely on statistical models and epidemiology research often encounters spatial and/or time series correlation between pollutants comparable to what we discussed for economic models. These associations confound efforts to isolate the pollutant responsible for specific health effects. Moreover, these joint effects can lead to the risk of double counting in the transfer process. This duplication can arise if one study identifies the culprit as pollutant A while another identifies pollutant B, and both are transferred. This issue was explicitly raised in EPA's early benefit analysis and is discussed in the health effects measures reported in the Prospective Analysis.

The ability of the damage-cost approach to separate effects and add them is an advantage in some respects, but a disadvantage in others. As pointed out by Randall [25] in a somewhat different context, adding up separate values is generally not consistent with the measure of total values when there are complementarities or income constraints. The question of income constraints may be particularly important given the large aggregate annual benefit estimates ($71 billion). A structural model such as ours provides a consistent way to treat diminishing marginal willingness to pay and income constraints. In principle, transferring unit values for morbidity could involve diminishing marginal willingness to pay per household, but usually they are assumed constant in real terms.[25]

Addressing questions about diminishing WTP for large changes in air pollution raises an inevitable question of what exactly are being designated as the ''outputs'' affected by policy. Air quality changes from the Clean Air Act Amendments over two decades can be expected to be large; yet if they are translated into health effects and considered in the context of total health

[24] Epidemiological estimates may also be biased downward if households systematically sort spatially based on their sensitivity to pollution, or engage in other averting behaviors. Our structural model is one way to explicitly address these types of sorting responses.

[25] In some cases real values are scaled to reflect differences in real income between the study and application areas. Smith et al. [31] propose a systematic method to reconcile diverse estimates and introduce such effects using a kind of structural meta-analysis of WTP estimates.

risks, their contributions to overall health changes may well be marginal. As a part of the baseline incidences of the health endpoints for the Prospective Analysis in 2010, they range from the largest example of 10.4% of cases of respiratory illness (attributed to nitrogen oxides and measured using hospital emergency room visits) to 0.003% for chest tightness, shortness of breath, or wheeze (attributed to sulfur oxides, and measured based on a chamber study using changes in the count of symptoms). Cases of premature death were 1% of the baseline incidences ([33] Table 5–3).

What is the relevant measure to be used in describing the consequences of reductions in ambient concentrations of pollutants and the associated improvement in air quality—these health effects, changes in visible range, and other amenity related effects (e.g. regional haze, materials and soiling effects, etc.)? Our framework and other revealed preference approaches estimate the value for the change in ambient concentration of an air pollutant directly and explicitly. The services being captured—whether actual health, perceived health, visibility as a service itself, or visibility as a proxy for other services—are unknown. The damage-cost approach makes these implicit services explicit. In so doing, however, it must make other assumptions. In particular, it must assume that expert epidemiologists and economists identify the ''right'' set of services. Typically, it also assumes that health risks from air pollution are equivalent to health risks posed by other sources, such as job risks and traffic accidents. This condition underlies the use of VSL measures from labor market studies to monetize the cases of premature death avoided by reductions in air pollution. It may also assume that risks to the sick or elderly are identical to risks to the young or healthy, or impose a linear relationship based on quality-adjusted life years.[26] Revealed preference models, such as ours, may incorporate all of these differences, but they do so in an implicit way. The explicit accounting used in damage function estimates allows cross-checks. For example, one can ask are the magnitudes of the physical impacts plausible?

Claims that the damage function approach offers the best methodology must rest ultimately on a more fundamental philosophical distinction with revealed preference methods. That is, these assertions amount to conclusions about whose judgments ultimately count in economic analysis. The principle of consumer sovereignty implies that people can generally be relied upon to know what is in their best interest. It is commonplace to assume in market demand studies people reveal these interests through their behavior. When experts know the details about various effects better than households, does it therefore follow that their estimates are more appropriate for benefit cost analysis? The question is much broader than non-market valuation. Cost-of-illness values for many market goods may be larger than any estimated consumer surplus. One might ask, for example, is the value of broccoli to be based on the demand for broccoli, or avoided cancer cases?

We suggest that, where they can be reliably estimated and transferred, revealed preference approaches have an important role to play in policy analysis. Efforts to present the two were a part of EPA's early benefit cost analyses. Attempts to reconcile or, at least compare them, have been eliminated from what seems to be current best practice. At a minimum, our argument implies this is a mistake. We have demonstrated how entire models capable of general equilibrium welfare analysis can be transferred, not simply the mean benefit measures they produce. The advantages

[26] Recent research by Alberini et al. [1] indicates it may be possible to recover measures of these VSLs for different groups with stated preference methods.

of this approach include the ability to account for these general equilibrium effects, as well as the resulting distributional effects. While the locational equilibrium framework is not yet ready for the types of full-scale policy analyses involved in our case study, our research confirms that there is merit in further refinement of the method.

6. Implications

The locational equilibrium framework provides a consistent basis for computing benefit measures for large and diverse changes in environmental conditions. The model identifies how diverse households (in terms of income and tastes for public goods) adjust to "the environmental cards they are dealt" by the combination of policy and nature. We have demonstrated it is possible to use an estimated locational equilibrium model with policy scale scenarios to describe how households would adjust to diverse environmental conditions.

As we noted at the outset, current practice uses both the damage function and unit benefit estimates in ways that assume households do nothing different as a result of the changes to environmental conditions. This assumption is unrealistic. Large policy interventions seem likely to generate different types of adjustments from those "embedded" in the reduced form ecologic or prospective cohort models used to describe health effects. Economic values change as well. The challenge is in consistently representing the market consequences of these adaptations. The locational equilibrium model relies on heterogeneity in tastes, income, and the levels of public goods to estimate preferences. When used to consider air pollution, we must assume that people recognize local air quality conditions, appreciate the consequences of that pollution, and use this information in their decisions about where to live. At a minimum, the current generation of locational equilibrium models has the ability to provide more detailed cross-checks for the conventional benefit estimates.

Acknowledgments

Thanks are due to Jim De Mocker of EPA, leader of the team of analysts responsible for the Prospective Analysis for assuring we were able to obtain the data necessary to complete this research, Allen Basala and Bryan Hubbell for helping us to understand them, to Leland Deck, James Neumann, and especially Kenneth Davidson for developing the data to match our modeling needs, to Maureen Cropper, Gilbert Metcalf, and Matt Kahn as well as the participants in NBER's Empirical Environmental Economics Conference and to Joseph Herriges and three anonymous referees for comments on several earlier drafts. Jaren Pope's excellent research assistance assured that the numerous Arcview issues required to link the geo-coded air quality data files for the EPA policy scenarios considered in this research would be compatible with our economic model. Susan Hinton assured the numerous drafts were prepared in a timely and consistent form. Partial support for this research was provided by the National Science Foundation NSF-S BR-98-08951, the US Environmental Protection Agency R828103, and the Alfred P. Sloan Foundation (for Sieg).

References

[1] A. Alberini, A. Krupnick, M. Cropper, N. Simon, J. Cook, The willingness to pay for mortality risk reductions: a comparison of the United States and Canada, Working Paper, October 2001.

[2] P. Bayer, R. McMillan, K. Rueben, An equilibrium model of sorting in an urban housing market: a study of the causes and consequences of residential segregation, Discussion Paper 03-01, Center for Economic Studies, Bureau of Census, 2003.

[3] S. Berry, Estimating discrete choice models of product differentiation, RAND J. Econ. 25 (1994) 242–262.

[4] S. Berry, J. Levinsohn, A. Pakes, Automobile prices in market equilibrium, Econometrica 63 (4) (1995) 841–890.

[5] N.E. Bockstael, K.E. McConnell, Public goods as characteristics of non-market commodities, Econ. J. 103 (420) (1993) 1244–1257.

[6] D. Burtraw, M. Cannon, Heterogeneity in costs and second best policies for environmental protection, Resources for the Future Discussion Paper No. 00-20, 2000.

[7] M.L. Cropper, L. Deck, N. Kishor, K. McConnell, Valuing product attributes using single market data: a comparison of hedonic and discrete choice approaches, Rev. Econ. Statist. 75 (2) (1993) 225–232.

[8] M.L. Cropper, L. Deck, K. McConnell, On the choice of functional forms for hedonic price functions, Rev. Econ. Statist. 70 (4) (1988) 668–675.

[9] W.H. Desvousges, F.R. Johnson, H.S. Banzhaf, Environmental Policy Analysis with Limited Information: Principles and Applications of the Transfer Method, Edward Elgar, Cheltenham, UK, 1998.

[10] B. Ellickson, Jurisdictional fragmentation and residential choice, Amer. Econ. Rev. Proc. 61 (2) (1971) 334–339.

[11] D. Epple, T. Romer, H. Sieg, Interjurisdictional sorting and majority rule: an empirical analysis, Econometrica 69 (6) (2001) 1437–1466.

[12] D. Epple, H. Sieg, Estimating equilibrium models of local jurisdictions, J. Polit. Econ. 107 (4) (1999) 645–681.

[13] J.A. Espinosa, V.K. Smith, Measuring the environmental consequences of trade policy: a nonmarket CGE analysis, Amer. J. Agr. Econ. 77 (1995) 772–777.

[14] L.H. Goulder, I.W.H. Parry, R.C. Williams III, D. Burtraw, The cost effectiveness of alternative instruments for environmental protection in a second best setting, J. Public Econ. 72 (3) (1999) 329–360.

[15] D.G. Hallstrom, V.K. Smith, Habitat protection and open space: a general equilibrium analysis of 'takings' and 'givings', Paper presented at AERE Workshop, Madison, WI, June 15–17, 2003.

[16] A.C. Harberger, Three basic postulates for applied welfare economics: an interpretive essay, J. Econ. Lit. 9 (1971) 785–797.

[17] M. Hazilla, R.J. Kopp, The social cost of environmental quality regulation: a general equilibrium analysis, J. Polit. Econ. 98 (4) (1990) 853–873.

[18] V.J. Hotz, R.A. Miller, Conditional choice probabilities and the estimation of dynamic models, Rev. Econ. Statist. 60 (1993) 497–530.

[19] M.F. Kokoski, V.K. Smith, A general equilibrium analysis of partial-equilibrium welfare measures: the case of climate change, Amer. Econ. Rev. 77 (1987) 331–341.

[20] R.C. Lind, Spatial equilibrium, the theory of rents, and the measurement of benefits from public programs, Quart. J. Econ. 89 (3) (1973) 188–207.

[21] Mathtech, Inc., Benefit and net benefit analysis of alternative national ambient air quality standards for particulate matter, Vol. I–IV, prepared for Office of Air Quality Planning and Standards, US Environmental Protection Agency, March 1983.

[22] R.B. Palmquist, Hedonic methods, in: J.B. Braden, C.D. Kolstad (Eds.), Measuring the Demand for Environmental Quality, North-Holland, Amsterdam, 1991.

[23] R.B. Palmquist, V.K. Smith, The use of hedonic property value techniques for policy and litigation, in: H. Folmer, T. Tietenberg (Eds.), International Yearbook of Environmental and Resource Economics 2002/2003, Edward Elgar, Cheltenham, UK, 2002.

[24] J. Poterba, Taxation and housing: old questions, new answers, Amer. Econ. Rev. Proc. 82 (2) (1992) 237–242.

[25] A. Randall, Total and nonuse values, in: J.B. Braden, C.D. Kolstad (Eds.), Measuring the Demand for Environmental Quality, North-Holland, Amsterdam, 1991.

[26] S. Scotchmer, Hedonic prices and cost/benefit analysis, J. Econ. Theory 37 (1985) 55–75.

[27] S. Scotchmer, The short-run and long-run benefits of environmental improvement, J. Public Econ. 30 (1986) 61–82.
[28] H. Sieg, V.K. Smith, H.S. Banzhaf, R. Walsh, Interjurisdictional housing prices in locational equilibrium, J. Urban Econ. 52 (2July) (2002) 131–153.
[29] H. Sieg, V.K. Smith, H.S. Banzhaf, R.P. Walsh, Estimating the general equilibrium benefits of large changes in spatially delineated public goods, Int. Econ. Rev., forthcoming.
[30] V.K. Smith, H.S. Banzhaf, A diagrammatic exposition of weak complementarity and the Willig condition, Amer. J. Agr. Econ., 2003, in press.
[31] V.K. Smith, G. Van Houtven, S.K. Pattanayak, Benefit transfer via preference calibration, Land Econ. 78 (2002) 132–152.
[32] D. Starrett, Land value capitalization in local public finance, J. Polit. Econ. 89 (2) (1981) 306–327.
[33] US Environmental Protection Agency, The benefits and costs of the Clean Air Act 1990 to 2010, Report to Congress, 1999.
[34] R.P. Walsh, Analyzing open space policies in a computable equilibrium model, Unpublished Ph.D. Dissertation, Duke University, 2002.
[35] R.D. Willig, Incremental consumer's surplus and hedonic price adjustment, J. Econ. Theory 17 (1978) 227–253.