Equilibrium Welfare Impacts of the 1990 Clean Air Act Amendments in the Los Angeles Area

by

Constant I. Tra

Working Paper 06-06

Department of Agricultural and Resource Economics

The University of Maryland, College Park

Equilibrium Welfare Impacts of the 1990 Clean Air Act Amendments in the Los Angeles Area[1]

Constant I. Tra
University of Nevada, Las Vegas
Constant.tra@unlv.edu

August 10, 2007

Abstract

This study develops a discrete choice locational equilibrium model to evaluate the benefits of the air quality improvements that occurred in the Los Angeles area following the 1990 Clean Air Act Amendments (CAAA). The discrete choice equilibrium approach accounts for the fact that air quality improvements brought about by the 1990 CAAA will change housing choices and prices. The study provides the first application of the discrete choice equilibrium framework (Anas, 1980, Bayer et al., 2005) to the valuation of large environmental changes. The study also provides new evidence for the distributional welfare impacts of the 1990 CAAA in the Los Angeles area. Households' location choices are modeled according to the random utility framework of McFadden (1973) and the differentiated product specification of Berry, Levinsohn and Pakes (1995). Findings suggest that the air quality improvements that occurred in the Los Angeles area between 1990 and 2000 provided an average equilibrium welfare benefit of $1,800 to households. In contrast, average benefits are $1,400 when equilibrium price effects are not accounted, demonstrating that ignoring equilibrium effects will likely underestimate the benefits of large environmental changes. In addition, we find that the equilibrium welfare impacts of the 1990 CAAA in the Los Angeles area varied significantly across income groups.

JEL Classification: H0, Q28, R13, R21

Keywords: benefit analysis, ozone improvement, locational equilibrium, discrete choice

[1] This paper is the result of my dissertation work at The University of Maryland. I thank my advisors Ted McConnell, Antonio Bento and Maureen Cropper for their guidance and support. I also thank Maximilian Auffhamer and Scott Lowe for their assistance in the early stages of this project. Thanks are also due to Robert Bernstein of the California Department of Education's Academic Accountability Unit, and Linda Nance of the California Criminal Justice Statistics Center. I am also grateful to V. Kerry Smith for his comments on an earlier version of this paper.

1 Introduction

Environmental regulations such as the Clean Air Act can lead to large air quality changes which cover wide areas and affect many residential communities. These types of widespread non-marginal improvements in air quality will have significant equilibrium welfare effects across local jurisdictions as households re-evaluate their residential location choices and equilibrium housing prices adjust. Traditional approaches to evaluating the impacts of air quality regulation have relied on direct welfare measures.[2] These welfare measures are recovered directly from the estimated preference function of consumers using either the hedonic framework (Rosen, 1974) or the discrete choice framework (McFadden, 1973, 1978). However, direct welfare measures do not explicitly account for the adjustments in housing prices which will occur when widespread non-marginal changes lead households to re-sort in the housing market. As a result, they will generally underestimate the full, i.e. equilibrium, welfare gains from regulations that result in widespread non-marginal improvements[3] in environmental amenities (Bartik, 1988, Palmquist, 1988).[4]

Recent studies by Sieg et al. (2004) and Smith et al. (2004) have shown that incorporating equilibrium adjustments can alter the estimates of welfare benefits from large environmental improvements. For instance, Sieg et al. (2004) find that the reductions in ozone levels during the five years following the implementation of the 1990 Clean Air Act Amendments led to equilibrium price increases ranging from 11 percent to 20 percent in the Los Angeles

[2] These welfare measures are often referred as partial equilibrium welfare measures.
[3] These are changes that are large enough to alter the stock of environmental quality in the market. As an example, consider the cleanup of all toxic waste sites in the Los Angeles metropolitan area.
[4] Equilibrium welfare measures are often referred as general equilibrium welfare measures.

metropolitan area. These price changes resulted in equilibrium welfare gains that were 13 percent higher than the direct benefits estimates that do not account for equilibrium adjustments.

This paper develops a discrete choice locational equilibrium model to evaluate the welfare benefits of the 1990 Clean Air Act Amendments (CAAA) to Los Angeles area households. The study makes two empirical contributions to public economics. First, the study provides the first application of the discrete choice equilibrium framework (Anas, 1980, 1982) to the valuation of large environmental changes. Households' location choices are modeled according to the random utility framework of McFadden (1978). The equilibrium model is closely related to the model of Bayer et al. (2005). This, more recent, discrete choice equilibrium model follows the differentiated product specification of Berry, Levinsohn and Pakes (1995) by incorporating unobserved attributes of residential locations in the household utility function. The discrete choice equilibrium framework provides an alternative to the framework proposed by Sieg et al. (2004) for evaluating the general equilibrium benefits of large environmental improvements. It allows for a richer and more realistic characterization of households' substitution patterns as well as preference heterogeneity.

Second, the paper provides new evidence for the distributional benefits of the CAAA in the Los Angeles area. Using the changes in ozone levels that occurred in the Los Angeles area between 1990 and 2000 we estimate average welfare benefits as well as the distribution of welfare benefits across income groups. Recently, Sieg et al. (2004) have provided estimates of the benefits of the CAAA based on the changes in ozone levels that took place between 1990 and 1995. With the availability of air quality monitoring data for the year 2000, we are able to evaluate the benefits of the CAAA from 1990 to 2000.

Little is known about the distribution of the benefits among households from the 1990 CAAA regulations. The only attempts at such an analysis have focused on the spatial distribution of welfare gains.[5] For instance, welfare gains in predominantly high-income neighborhoods are compared with those in low-income neighborhoods. This approach, however, fails to capture the distribution of welfare gains and losses across household characteristics such as income and race. It only provides a comparison of the welfare gains across neighborhoods.

Household preferences are estimated using a dataset which includes households and housing units from the 1990 Census Public Use Microdata Sample (PUMS), annual ozone summaries from the California Air Resource Board, school performance data from the California Department of Education and crime indices from the California Criminal Justice Statistics Center. Households' residential location choices are characterized by a discrete choice model in which equilibrium conditions are enforced. The model captures the heterogeneity of household preferences for location amenities by incorporating observed household characteristics in the utility function. Observed household characteristics include household income, household size, employment location and educational attainment of the household head.

Estimation of the equilibrium welfare impacts incorporates price adjustments that result from the fact that households alter their residential location choice after the changes in air quality throughout the Los Angeles area. Computation of the equilibrium adjustments is obtained via simulation. Using 1990 as a benchmark we simulate market clearing prices and household choices for the counterfactual locational equilibrium that would have resulted in 1990 if air quality levels were identical to those observed in 2000, while all other housing attributes

[5] See for example Shadbegian et al. (2004). Smith et al. (2004) investigate the distributional impacts of the 1990 CAAA using the projected air quality changes, in the Los Angeles area for the year 2000, from the EPA's 1999 prospective study. However, the actual air quality changes between 1990 and 2000 significantly differ from the EPA projections.

remained at their 1990 levels. The counterfactual equilibrium only captures the changes in air quality that occurred in the Los Angeles area between 1990 and 2000. Other factors characterizing the Los Angeles area housing market, such as population, household income and housing supply, are assumed fixed in the simulation.

The empirical analysis focuses on the four counties of the Los Angeles area which makeup the South Coast Air Quality Management District. This area experienced significant improvements in air quality during the decade that followed the implementation of the 1990 CAAA. The results suggest that the reductions in ozone concentrations across Los Angeles, Orange, Riverside and San Bernardino counties, provided an average equilibrium benefit of $1,800 to households. In contrast, average benefits are $1,400 when equilibrium adjustments are not accounted, demonstrating that ignoring equilibrium effects will likely underestimate the benefits of large environmental changes. We find that the equilibrium welfare impacts of the 1990 CAAA in the Los Angeles area varied significantly across income groups. Households in the highest income quartile experienced equilibrium benefits of approximately $3,600 as compared to only $400 for households in the lowest-income quartile. We also find that ignoring equilibrium adjustments can significantly alter the distribution of relative welfare gains (i.e. welfare gains as a proportion of household income). Indeed, welfare impacts that do not account for equilibrium effects suggest that high-income households have larger relative welfare gains compared to low-income households. However, when accounting for equilibrium adjustments, we find that the distribution of relative welfare gains from the 1990 CAAA is fairly even across income groups.

The remainder of this paper is organized as follows. Section 2 provides some background information and reviews the current body of literature on the valuation of housing amenities.

Section 3 characterizes the locational equilibrium model. Section 4 describes the various datasets used to estimate the household utility function. Section 5 outlines the estimation of the household preference parameters. Section 6 discusses the simulation and the welfare results. Section 7 concludes the analysis.

2 Background and Literature Review

2.1 Background

The 1990 Clean Air Act Amendments [6]

The Clean Air Act Amendments (CAAA) of 1990 addressed three major environmental issues in the United States: acid rain, urban air pollution and toxic air emissions. Title I established new provisions for the attainment and maintenance of the National Ambient Air Quality Standards (NAAQS). It is intended to address the urban air pollution problems arising from ground-level ozone, carbon monoxide and particulate matter (PM-10). Areas for which ambient levels of these pollutants were above the target levels were designated as non-attainment areas by EPA. Non-attainment counties for ozone were classified into five categories (marginal, moderate, serious, severe and extreme). These areas were then required to implement control measures that vary with the severity of their non-attainment status. For carbon monoxide and particulate matter, areas that did not meet the federal health standards were classified into either moderate or serious non-attainment status. Areas exceeding carbon monoxide standards were required to introduce oxygenated fuels programs and/or implement enhanced emission inspections. Depending on the severity of their status, particulate matter non-attainment counties were either required to

[6] Based on U.S. Environmental Protection Agency (2006a).

implement reasonably available control measures (RACM) or best available control measures (BACM).

Air Quality Standards for Ground-level Ozone [7]

Under the Clean Air Act, EPA is required to set National Ambient Air Quality Standards (NAAQS) for pollutants that are considered to be harmful to public health and the environment. Currently, two standards are used to regulate ozone levels in the U.S. The national 1-hour standard for ozone, set at 0.12 parts per million (ppm) by volume, was established in 1979. It is achieved when the average number of days per calendar year with maximum hourly concentrations above 0.12 ppm does not exceed 1. In 1996, EPA established a new national 8-hour ozone standard which was set at 0.08 ppm by volume. This standard is attained when the three-year average of the fourth highest daily maximum 8-hour ozone concentration measured at each ozone monitor within an area is less than 0.08 ppm. In June of 2005, the 1-hour ozone standard was revoked in all areas and replaced by the 8-hour standard, except in the fourteen 8-hour ozone non-attainment areas that were part of EPA's Early Action Compacts[8] program.

In addition to setting the NAAQS, EPA designates areas as either non-attainment, attainment or unclassified. The designation process plays an important part in the implementation of air pollution control measures by states and local governments. Currently, an area is designated as non-attainment if it violates the national 8-hour ozone standard over a three-year period. An area will be designated as attainment if it has air quality monitoring data showing that the area has not violated the ozone standard over a period of three years. Areas are designated as unclassified if there are not enough data to determine ozone levels.

[7] Based on U.S. Environmental Protection Agency (2006b).

[8] Early Action Compacts give local communities the flexibility to develop their own approach to meeting the 8-hour ozone standard, provided the communities control emissions from local sources earlier than the Clean Air Act would otherwise require.

Air Quality Improvements in the Los Angeles Area

The South Coast Air Quality Management District (AQMD) is the main regulatory body for air pollution in the Los Angeles area. It encompasses Orange County and the urban areas of Los Angeles, Riverside and San Bernardino County. The area is the most densely populated urban center of the state of California and is home to over 16 million people. The South Coast Air Quality Management District has historically exhibited some of the worst ambient levels of air quality in the nation (U.S. EPA, 2006c). Every three years AQMD develops an air quality management plan which identifies implementation measures designed to bring the area in compliance of state and federal air quality standards.

Figure 1 provides maps of ozone concentrations in 1990 and 2000 for the four counties which makeup the South Coast AQMD. The 1990 map shows a wide variation in ozone levels across the area. Specifically, ozone concentrations were lowest in the coastal areas of Los Angeles and Orange County. Average 1-hour ground-level ozone concentrations, in those areas, were below the federal 1-hour standard (0.12 ppm). On the other hand, the areas east of the San Bernardino Mountains and south of the San Gabriel Mountains exhibited the highest ozone concentrations in 1990. Average 1-hour ground-level ozone concentrations in these areas ranged from 0.185 ppm to as high as 0.225 ppm.

The South Coast AQMD counties experienced significant reductions in ozone concentrations between 1990 and 2000. Table 1 reports average ozone concentrations from monitoring stations across the area. The average 1-hour ground-level ozone reading in 2000 was roughly 0.10 ppm, compared to 0.14 ppm in 1990. In addition, the number of days exceeding the federal 1-hour standard (0.12 ppm) significantly decreased between 1990 and 2000. The average number of recorded exceedences across the area was about 3.5 days in 2000, compared to nearly 36 days in

1990. Figure 1 shows that the ozone reductions were highest in areas with the worst ground-level ozone concentrations in 1990. Average ozone concentrations fell by nearly 62 percent at monitoring stations with a recorded 1990 ozone level above the federal 1-hour standard (0.12 ppm). On the other hand, monitors with a recorded 1990 ozone level below the federal 1-hour standard experienced an average reduction of only 28 percent.

2.2 Approaches to Valuing Amenity Changes

The type of empirical approach required to value amenity changes ultimately depends upon the question of interest. Researchers are generally interested in estimating the marginal value from amenity changes. This is obtained by estimating the hedonic price function for the relevant housing market and taking the gradient with respect to the amenity of interest. Smith and Huang (1995) provide an extensive survey of the applications of this approach to air quality valuation. Sometimes a researcher might want to estimate the willingness to pay (WTP) of households for a non-marginal amenity change at their residential location. This requires estimating the demand or WTP function for the amenity. The welfare estimation can be implemented via Rosen's (1974) two-stage hedonic approach or McFadden's (1978) discrete choice approach. Palmquist (2006) provides a recent review of hedonic and discrete choice demand approaches to valuation.

The hedonic and discrete choice demand estimations do not, however, allow the researcher to evaluate the welfare impact resulting from large amenity changes. In fact, these methods will underestimate the welfare gains from large amenity changes because they do not account for the fact that households change their locations. Evaluating the welfare impact of large amenity changes is generally a more complex task as one needs a model that explicitly incorporates changes in household location choices.[9]

In contrast to the hedonic and discrete choice demand estimation, which assume that non-marginal amenity changes do not affect household location choices, locational equilibrium models are able to incorporate price adjustments that result from the re-sorting of households across housing locations in response to a policy change. These models use estimated household preferences to simulate a counterfactual equilibrium outcome for a policy change.

Sieg, Smith, Banzhaf and Walsh (2004) provide the first empirical analysis of the equilibrium welfare impacts from non-marginal environmental improvements.[10] The study develops a locational equilibrium model based upon Epple and Sieg's (1999) equilibrium framework. Households, in the Sieg et al. model, select housing locations among a finite set of differentiated communities. The set of communities is characterized by 91 school districts. Communities differ in their provision of local public goods (including air quality) and housing. Conditional on their community choice, households select housing as a continuous, homogeneous good. Households' preferences vary with respect to income and a taste parameter. The locational equilibrium is defined in terms of three conditions: boundary indifference, income stratification and ascending bundles. These properties are used to estimate parameters of the household's utility function, which are then used to simulate alternative equilibrium outcomes for changes in ozone concentrations at the school district level.

Sieg et al. (2004) apply this framework to investigate the welfare benefits of the 1990 CAAA in the Los Angeles area. They find that equilibrium benefits that account for adjustments in

[9] There is one instance when evaluating the welfare impact of a large amenity change is a simple task. This occurs when the amenity change is localized, i.e., confined to a small geographic area. In this case it can be shown (see Bartik, 1988) that the welfare impact will equal the sum of the changes in housing prices within the affected area.
[10] See also Smith et al. (2004) and Walsh (2003) for other applications related to the Sieg et al. (2004) approach.

housing prices differ substantially from direct benefit estimates. The average equilibrium welfare gain from the reductions in ozone concentrations, which occurred between 1990 and 1995 in the Los Angeles area, was estimated at $1,371. This compares with the average direct benefit of $1,210. In addition, the study finds a significant amount of heterogeneity in welfare gains across counties. Equilibrium benefits were found to be highest in Los Angeles County ($1,556) and lowest in San Bernardino County ($367). The study also finds considerable variation in benefits across school districts, within each county. For example, the equilibrium benefits in Los Angeles County ranged from $486 in the Compton Unified School District to $9,000 in the Beverly Hills School District.

In a subsequent study, Smith et al. (2004) evaluated the benefits of the 1990 CAAA in the Los Angeles area for 2000 and 2010. Using the EPA's projected changes in ozone levels for 2000 and 2010 together with the estimated household preferences from Sieg et al. (2004), the study measures the equilibrium WTP for the policy scenarios developed for EPA's prospective study (EPA, 1999) as they relate to the households of the Los Angeles area. The study also investigates the distribution of equilibrium benefits across income groups. They present the benefits associated with the 25^{th}, 50^{th} and 75^{th} income percentiles for selected school districts in the Los Angeles area. The estimated equilibrium welfare estimates vary significantly across the household income distribution. The distribution of the welfare estimates also varies across school districts. In the lowest-income community, San Juacinto Unified School District, the welfare estimates are -$59 annually for the 25^{th} income percentile as compared to -$28 for the 75^{th} percentile. The welfare estimates in Beverly Hills School District, the highest-income community, are $3899 for the 25^{th} income percentile as compared to $7406 for the 75^{th} percentile.

Sieg et al. (2004) provide a major contribution to the valuation of large widespread changes in environmental amenities. The study provides the first explicit characterization of the equilibrium impact of non-marginal amenity changes on household choices and housing prices. In addition, the Sieg et al. model has the advantage of being simple to implement and computationally tractable even with a large number of housing locations. However, the specification of household preferences, that is needed to ensure that the necessary conditions for the equilibrium are met, gives rise to some limitations.

First, the characterization of preferences for amenities leads to somewhat restrictive patterns of substitution across locations. This is because location amenities enter the household's indirect utility function through a single public good index. As a result, households are forced to have the same ranking of communities in the amenity space. This vertical differentiation of communities simplifies the estimation of preference parameters and the computation of the locational equilibrium. However, one would generally expect households to have different relative preferences for community-specific amenities such as air quality, education and crime. For instance, other things equal, one would expect that households with children enrolled in a secondary public school will have higher preferences for communities with good secondary public schools.

A second limitation of the Sieg et al. (2004) model relates to the characterization of the heterogeneity in households' preference for location amenities such as air quality, school quality and crime. Heterogeneity in households' preferences for the public good index is characterized by a single taste parameter whose marginal distribution is assumed normal. Hence a household's marginal valuation for a given community amenity is only a function of the household's income and does not depend on other household characteristics. Households' preferences for community-specific attributes are, however, likely to vary across other household characteristics such as household size, the presence of children and educational attainment. For instance, highly educated household are likely to have a higher marginal valuation for school quality. As a result, a preference specification which incorporates an interaction between neighborhood school quality and household educational attainment would allow the model to better fit the data. In addition, when investigating welfare gains from an amenity change, a researcher is able to provide an analysis of the distributional impacts across household characteristics other than income. For instance, one may want investigate the differential impact of an improvement in air quality on minority households.

An alternative to the Sieg et al. (2004) equilibrium framework is the discrete choice equilibrium framework. This is the equilibrium approach adopted here. Anas (1980, 1982) developed a theory of locational housing market equilibrium based on the discrete housing choice model of McFadden (1978). In recent years this framework has been extended to incorporate advances in urban economics and empirical industrial organization. One such model was proposed by Bayer and Timmins (2005). Their model incorporates endogenous social interaction effects as well as unobserved location attributes.

The discrete choice equilibrium approach provides for a richer characterization of preference heterogeneity and more general patterns of substitution. The discrete-choice modeling of the household location allows community-specific amenities to enter directly the utility function. This provides for more general substitution patterns across communities. In addition, the researcher can characterize the observed heterogeneity in households' tastes for location amenities by incorporating interactions of household characteristics and location amenities into the utility function. This would allow the researcher to evaluate the impact of a policy change on various socio-economic subgroups of the household population. The main limitation of the discrete choice equilibrium approach is that its implementation requires significant computational work.

To date, discrete choice equilibrium models have been mostly used to analyze urban and transportation policy changes. Anas (1982) evaluates the impact of public transportation projects proposed for the Chicago area in the early 1980s. Bayer et al. (2005) use an equilibrium model similar to the Bayer and Timmins (2005) model to investigate the impact of an increase in income inequality in the San Francisco bay area. Timmins (2007) uses a similar equilibrium model to evaluate the welfare costs of rainfall changes in Brazil using labor market data. The equilibrium model in this paper is based on the specification of Bayer et al. (2005).

3 A Locational Equilibrium Model for the Los Angeles Area

This section develops the discrete choice equilibrium model used to evaluate the welfare impacts of the 1990 Clean Air Act amendments in the Los Angeles area. We model households' location decisions according to the framework of Bayer et al. (2005). The characterization of the locational equilibrium follows Anas (1982).

3.1 Modeling Households' Location Choice

Households are assumed to choose their residential location h from a discrete set of housing types (H). A housing type is defined as a collection of houses with identical observed characteristics and located within the same neighborhood. The utility that a household i derives from a residential location h is given by:

$$v_{ih} = \alpha \log(y^i - p_h) + \gamma d_{ih} + \sum_k x_{hk} \beta_{ik} + \xi_h + \varepsilon_{ih}\,, \tag{3.1}$$

where y^i represents household i's monthly income and p_h is the monthly rental price of house h. d_{ih} is a dummy variable which equals 1 if the residential location is within the household's employment zone. It is intended to capture the household's preference for housing locations that are closer to its workplace. The k^th element of the vector of observed attributes for residential location h is given by x_{hk}. These are the housing and neighborhood attributes that are present in the researcher's data. Housing characteristics include bedrooms, age, dwelling type and tenure status. Neighborhood characteristics include ozone concentration, 8^th grade math score, crime index, elevation, proximity to the Pacific coastline, housing density and proportion of Hispanics. Other attributes of the residential location that are observed by the household but not observed in the data enter the household's utility via the location-specific error term, ξ_h. This term will capture the household's average valuation of the unobserved attributes.[11] The last term, ε_{ih}, is a mean-zero stochastic error which captures the unobserved taste heterogeneity among households.

Each household chooses the residential location which provides it with the highest utility. The household's indirect utility derived from this maximization problem is given by:

$$V_{ih} = \underset{h \in H}{\text{Max}}\;\; \alpha \log(y^i - p_h) + \gamma d_{ih} + \sum_k x_{hk} \beta_{ik} + \xi_h + \varepsilon_{ih}\,, \tag{3.1a}$$

[11] As in Bayer et al. (2005) and Berry et al. (1995), the specification of the indirect utility assumes that households have the same valuation for the unobserved attributes. Hence, we are not able to identify heterogeneous preferences for unobserved location attributes. Athey and Imbens (2007) use Bayesian methods to estimate a random utility model which incorporates individual-specific tastes for the unobserved attributes.

where α, γ and β_i are parameters of the household's preference function. α characterizes the household's marginal utility of the log of income[12] while β_{ik} captures the household's taste for location attribute k. The parameter γ characterizes the household's disutility for commuting to work. We explicitly account for the heterogeneity in households' preferences for location characteristics by allowing the taste parameters to vary systematically across households. The specification of the heterogeneous taste parameters uses interactions between location characteristics and observed characteristics of households. These observed household characteristics include household income, household size, the presence of children under the age of 18 and whether the household head is college educated. The functional form of the household's taste for an attribute k is given by:

$$\beta_{ik} = \beta_{0k} + \sum_r z_{ir}\beta_{1kr}\,, \tag{3.2}$$

where z_{ir} represents the r^{th} characteristic of household i. The first term captures the component of the household's taste for the attribute k which is common across all households.[13] The second term is intended to capture systematic differences in tastes which can be attributed to the household's observed characteristics.

The final form of the indirect utility function is obtained by substituting equation (3.2) into equation (3.1a) for the chosen location. It is given by:

---

[12] So that the marginal utility of income is given by α / (y_i - p_h).

[13] When the household characteristics (z_{ir}) are demeaned, this term will equal the mean taste parameter across households.

$$V_{ih} = \delta_h + \alpha \log(y^i - p_h) + \gamma d_{ih} + \sum_{kr} x_{hk} z_{ir} \beta_{1kr} + \varepsilon_{ih} \,, \tag{3.3}$$

where,

$$\delta_h = \sum_k x_{hk} \beta_{0k} + \xi_h \,. \tag{3.4}$$

Equation (3.3) outlines the two main components of the household's valuation of its chosen location. The first component, represented by the constant term (δ_h), captures households' common valuation of location attributes. This valuation is shared by households regardless of their characteristics. For instance, all else equal, households would prefer a house with more bedrooms, less pollution, better schools, less crime, etc. This common valuation represents the base utility that households derive from the residential location *h*.[14] The second component captures households' individual valuation of the location attributes. These individual valuations are assumed to arise from differences in the observed characteristics of households. For instance, all other things equal, households of larger size are likely to choose houses with more bedrooms.

Bayer et al. (2005) suggest a two-stage approach to estimate the parameters of the household location choice model in equation (3.3). In the first stage, one would recover the household-specific taste parameters (α, γ, β_1) and the location-specific constants (δ_h). This stage can be implemented by maximum likelihood estimation. Because of the large number of housing types the alternative constants are estimated using the contraction mapping proposed by Berry et al. (1995). The details of the estimation are provided in Section 5. The second stage then estimates

[14] When the household characteristics (z_{ir}) in equation (3.2) are demeaned, this base utility will also represent the mean utility provided by the residential location *h*.

the mean taste parameters (β_{0k}) from the regression specification provided by equation (3.4) using the location constants estimated in the first stage.

The household utility in equation (3.3) closely resembles the utility specification in Bayer et al. (2005). However, there are two differences between our specification and that of Bayer et al. (2005). One difference arises from the characterization of the non-housing good. We characterize the household's consumption of the composite non-housing good using the term $log(y_i - p_h)$. This allows the model to capture income effects that are present in the household's choice problem. It also allows us to derive Hicksian welfare measures that are consistent with the household's utility maximization problem. In the Bayer et al. model the indirect utility does not incorporate the composite non-housing good. The household income enters the utility as a linear interaction with location attributes, and the housing price enters the utility linearly as an attribute of the residential location.

The second difference between our model and the model used by Bayer et al. (2005) is that we do not incorporate endogenous social interaction effects. Social interaction effects emerge from the fact that households may care about the average socioeconomic characteristics of their neighborhoods. These social interaction effects are likely to be endogenously determined in the sorting equilibrium when households have heterogeneous preferences. This is because the average socioeconomic makeup of neighborhoods changes each time households resort. In our utility function the social interaction effect arises from households' homogeneous tastes for the proportion of Hispanics in their neighborhood. As a result the social interaction effect is exogenous.

Our specification of the household's indirect utility differs fundamentally from Sieg et al. (2004). Sieg et al. specify the indirect utility of a household residing in a community j as:

$$V_{ij} = [\alpha g_j{}^{\rho} + h(y_i, p_j)^{\rho}]^{1/\rho}, \tag{3.5}$$

where g_j is the public good index for community j and $h(\bullet)$ is a non-linear function characterizing the household's expenditures on housing. y_i represents the household's income while α_i is a parameter characterizing the heterogeneity of the household's valuation for the public good index. p_j represents the housing price index for community j.

Two main distinctions arise between our equilibrium model and the model used by Sieg et al. (2004). First, according to the Sieg et al. specification, households value community amenities through the single public good index g. As a result, households will have the same preference ordering of communities in the amenity space. This type of preference structure generates substitution patterns that can be restrictive since households are forced to have the same ranking of communities in the amenity space. In our specification, substitution patterns are determined by the interaction of household characteristics and location attributes. Hence, households will have different relative preferences for community-specific amenities such as air quality, education and crime.

Second, the interaction of household characteristics and location attributes also provide a richer characterization of the heterogeneity in household preferences for location amenities. The taste heterogeneity with respect to the community air quality level is captured by interaction with household income. Heterogeneity in preferences for school quality is captured via interaction with the household's educational attainment. This approach differs from the Sieg et al. model where heterogeneity in preferences for amenities is characterized by the single unobserved taste parameter, α.

3.2 Characterizing the Locational Equilibrium

We now turn to the characterization of the locational equilibrium for the housing market. We first derive the predicted demand for each housing type. The demand side of the market is made of N heterogeneous households. The supply side of the housing market comprises N occupied housing units classified into H housing types. The supply of each housing type h is defined as the measure of housing units of type h in the study area and is assumed fixed. The locational equilibrium defines a set of market clearing prices $\{p_h\}$ and household choice probabilities $\{P_{ih}\}$.

Characterizing the Housing Demand

We will assume that the idiosyncratic error component ε_{ih} is identically and independently distributed and has a Type I Extreme Value (EV) distribution. Given this assumption, the probability that a household chooses a residential location h is defined by:

$$P_{ih}(p, z_i, x) = \Pr[V_{ih} > V_{il,\forall l \neq h}] = \frac{\exp[\delta_h + \alpha \log(y_i - p_h) + \gamma d_{ih} + \sum_{kr} x_{hk} z_{ir} \beta_{1kr}]}{\sum_{m=1}^{H} \exp[\delta_m + \alpha \log(y_i - p_m) + \gamma d_{im} + \sum_{kr} x_{mk} z_{ir} \beta_{1kr}]}, \tag{3.6}$$

The predicted aggregate demand for housing type h is obtained by summing the choice probabilities (P_{ih}) over the household population.

$$d_h(p) = \sum_i P_{ih}(p, z_i, x), \tag{3.7}$$

where p is a vector of housing prices, z_i is a vector of housing characteristics and x is a matrix of location attributes whose columns are x_h.

Equation (3.6) characterizes a multinomial logit (MNL) choice structure. An implication of the MNL choice structure is the independence from irrelevant alternatives (IIA) property, which has been the subject of much criticism in the discrete choice literature. A direct consequence of the IIA property is that, for a given household, the ratio of the choice probabilities for any two alternatives is independent of the household's systematic valuation of the remaining other alternatives in the household's choice set. It should be noted that while IIA is a property of the individual household choice probabilities in our model, it is not a property of the housing demands. This can be easily seen by looking at the ratio of the predicted demands for housing alternatives *k* and *l*:

$$\frac{d_h(p)}{d_l(p)} = \frac{\sum_i P_{ih}}{\sum_i P_{il}} = \frac{\sum_i 1/\left(1+\sum_{m\neq h}\exp[V_{im}]\right)}{\sum_i 1/\left(1+\sum_{m\neq l}\exp[V_{il}]\right)}. \tag{3.8}$$

It is clear that the ratio in equation (3.8) is not independent of the remaining housing alternatives in the choice set. The only instance when this ratio can be independent of the remaining alternatives is when households have identical characteristics. In this case the ratio equals one. Hence, the inclusion of household characteristics in the indirect utility function ensures that the housing demands derived from the model will exhibit realistic substitution patterns.

Defining the Locational Equilibrium

The supply of housing units of type h, s_h, is assumed fixed and is given by the number of housing units of type h in the data. The locational equilibrium is such that the demand for each housing type equals its supply. It is characterized by a vector of H housing prices p and a set of NH

household location choice probabilities $\{P_{ih}\}$. More specifically, the vector of market-clearing prices p is defined by:

$$d_h(p) = s_h \qquad\qquad h = 1,..., H. \tag{3.9}$$

Equation (3.9) defines a system of H equations in H variables. Anas (1982) shows that a unique vector of market-clearing prices exists when the household location choice probabilities P_{ih} are strictly decreasing in the housing price p_h. This occurs when the estimate for the parameter α is positive.

4 Data Sources

The focus of this study is on the four counties that make up the South Coast Air Quality Management District (AQMD): Los Angeles County, Orange County, Riverside County and San Bernardino County. We estimate the parameters of the household preference function defined by equations (3.3) and (3.4) using a cross-section of 1990 microdata which includes household characteristics, housing characteristics, neighborhood air quality, neighborhood school quality, neighborhood crime rate, neighborhood racial composition, neighborhood housing density, neighborhood elevation and proximity of the neighborhood to the Pacific coastline. The remainder of this section describes the housing and air quality data. A description of the remaining data is provided in the appendix.

4.1 Household and Housing Characteristics

Households and housing characteristics are obtained from the 1990 Census 5 percent Public Use Microdata Sample (PUMS).[15] These are records containing a 5 percent sample of all housing units in the United States. The PUMS records provide an extensive description of the housing stock and the households in the occupied dwelling units. The PUMS are extracts from the actual decennial Census long-form questionnaire, which are taken in a way that protects the confidentiality of households. However, unlike the confidential long-form files, which identify each household's Census block (an area of approximately 100 people), the 5 percent PUMS sample only identifies the location of households in a Public Use Microdata Area (PUMA), which is a Census geographic area containing approximately 100,000 people. The PUMS also identify the employment location of household members by their workplace PUMA.

The 1990 PUMS sample for the four counties in the study area comprises 224,565 occupied housing units. The original household sample consists of the 224,565 households that occupy those housing units. Our analysis focuses on the households occupying single and multi-family dwelling units. Mobile homes and group quarters are excluded from the sample. In addition, we restrict our sample to households that have a monthly income of at least five hundred 1990 dollars. Finally, we dropped the observations where the household's reported monthly income was less than the monthly rental price of the housing unit. The final sample, which is used to represent the population of households and housing units in this study, consists of approximately 171,000 observations.

Sieg et al. (2004) estimate household preference parameters using housing transactions from 1989 to 1991 in Los Angeles, Riverside, Orange and Ventura County. These data identify the

[15] These data are publicly available from the U.S. Census Bureau (www.Census.gov), or at www.ipums.umn.edu/usa/vars.html.

Census tract in which a housing unit is located. Sieg et al. characterize residential communities using 1990 school district boundaries. Housing transactions data provide a more comprehensive set of housing characteristics than the Census long form. However, these data do not provide information on the households occupying the houses. As a result they do not allow one to estimate richer preference specifications, such as those used Bayer et al. (2005), where preferences for location amenities vary across household characteristics.

Table 2 provides descriptive statistics of the household and housing characteristics in our 1990 PUMS sample. The microdata sample comprises 171,000 observations describing households and their occupied housing units. The vast majority (nearly 70 percent) of the households in the study area reside in Los Angeles County. Orange County has the second most household population (17 percent), followed by San Bernardino County (10 percent) and Riverside County (3 percent).

The average number of bedrooms for the houses in the study area is 2.25. We follow the approach of Bayer et al. (2005) to compute an imputed monthly rental housing price across owner-occupied and rental units. A detailed description of the method is provided in the appendix.[16] The mean monthly rental housing price is \$749. Monthly housing prices are highest in Orange County (\$956) and lowest in San Bernardino County (\$707). Half of the housing units in the study area are owner-occupied. Riverside and San Bernardino County have the largest owner-occupied housing shares (0.63). Overall, the housing stock is quite young. Nineteen percent of the houses were built after 1980; 37 percent were built in the 1960s and 70s.

[16] We construct a single price vector for owned and rental housing units by estimating a hedonic price regression for each of the three metropolitan statistical areas in the PUMS sample (Los Angeles-Long Beach, Orange County and Riverside-San Bernardino). The regressions provide an estimate of the average ratio of housing values to monthly rents in each metropolitan statistical area. The average ratio for the study area is 316.1. The average ratios are then used to convert housing values to their corresponding rental rates.

A household's preference for housing locations that are closer to its workplace is captured by a dummy variable which equals 1 if a residential location is within the household's employment zone. The household's employment zone is defined as the PUMA of the household head's workplace. Other studies (see e.g., Bayer et al., 2005 and Takeuchi et al., 2005) have instead used the distance to the householder's employment location. However, in the PUMS data, the householder's employment location is given by the workplace PUMA. Hence the distance to the household's employment location cannot be calculated. Because the workplace PUMA is a relatively large geographic area we prefer using a dummy variable for whether the residential location is within the workplace PUMA, instead of the distance from the residential location to the workplace PUMA. The latter turns out to be a noisier measure. Roughly half of the households in the study area choose housing units which are located within their employment zone.

The lower half of Table 2 provides a summary of means for the household characteristics that enter the model. The average monthly household income in the sample is $4,098. Orange County has the highest average monthly income ($4,945), whereas Riverside County has the lowest average ($3,860). The racial profile of the household is given by the race of the household head. The study area comprises 8 percent non-Hispanic Asian and 9 percent non-Hispanic Black households. Fifty-eight percent of the households in the study area are non-Hispanic Whites. Households of Hispanic origin make up 23 percent of the population. The share of Hispanic households is highest in Los Angeles County (26 percent) and lowest in Orange County (15 percent). The educational attainment of the householder is captured by a binary variable indicating whether the household is college educated. Thirty-five percent of households in the study area are headed by a college graduate.

4.2 Neighborhood Variables

Defining Neighborhoods

Table 3 reports average values for the neighborhood attributes used in the model. We use the 1990 Census PUMA boundaries to characterize neighborhood geography. This is because the PUMS identifies the geographic location of a dwelling unit as the Census PUMA. A Census PUMA is a geographic area containing approximately 100,000 individuals. Sieg et al. (2004) characterize residential communities using 1990 school district boundaries. They were able to do so because housing transactions identify the Census tract as well as the school district for each housing unit. Because they had access to the 1990 Census long-form files, Bayer et al. (2005) were able to use Census block boundaries to characterize neighborhoods. The Census block is a geographic area of approximately 100 individuals.

The study area comprised a total 87 PUMAs in 1990. This compares with approximately 150 school districts and 2,400 Census tracts. The average PUMA in 1990 had approximately 3,000 housing units. To reduce measurement errors in characterizing neighborhood attributes, the estimation only uses PUMAs whose boundaries are mutually exclusive. PUMAs that are enveloped by other PUMAs are excluded from the sample. This reduces the number of PUMAs to 79.

PUMAs are relatively large geographic units compared to Census tracts or school districts. However, for the main attributes used in the estimation, the variation within PUMAs is small compared to the variation across PUMAs. Table 4 shows within and between PUMA standard deviations for selected characteristics. For math score, ozone and PM-10 values, the variation across PUMAs is nearly five times larger than the within PUMA variation. The difference is smaller for the crime measure. The standard deviation of crime values across PUMAs is 20 percent higher the mean standard deviation within PUMAs. We therefore conclude that the PUMA boundaries provide a good characterization of neighborhood school quality, crime and air quality.

Air Quality Data

The air quality data used in this study were obtained from the California Air Resources Board (CARB). CARB provides California ambient air quality data for criteria and toxic pollutants from 1980 through 2002. The data include hourly and daily values as well as annual summaries collected from a large network of monitors dispersed throughout the state of California. Annual averages for 1990, 1995 and 2000, are obtained for two major primary criteria pollutants: ozone and particulate matter (PM-10). These pollutants have been shown to have a significant impact on housing prices (Sieg et al., 2004). Ozone is measured as the average of the top-30 daily maximum readings at a monitor, while particulate matter (PM-10) is measured by the annual geometric mean.

Table 1 provides descriptive statistics of the monitor air quality data in the study area. Average ozone concentrations in 1990 were highest in Los Angeles County and lowest in Orange County. Ozone concentrations fell by nearly 40 percent between 1990 and 2000, with the largest reductions recorded in the worst areas. Monitor readings tend to be strongly correlated across pollutants. Table 5 shows the correlation between ozone, PM-10, nitrogen oxide (NOx) and sulfur dioxide (SO_2). The correlation coefficient for ozone and PM-10 at monitor locations measuring both pollutants is 0.44. Ozone and PM-10 levels are also strongly correlated with secondary pollutants such as nitrogen oxide and sulfur dioxide. The correlation coefficient between ozone and NOx is 0.47; for ozone and SO_2 it is -0.56.

The study area had a total of 50 active monitors measuring ozone between 1989 and 1991. This compared with 18 monitors measuring PM-10 concentrations. Interpolation techniques are used to determine neighborhood air pollution levels. We use two approaches to determine neighborhood air pollution levels. The first approach assigns to each PUMA the centered three-year average of readings from the closest monitor. If more than one monitor falls within a PUMA, the PUMA is assigned the average from these monitors. Sieg et al. (2004) used a similar approach to assign air quality levels to each house in their sample. They then approximate the neighborhood air quality level using the averages for the houses sold in each school district. One potential issue with this approach is that it may assign the same monitor readings to a collection of neighborhoods, regardless of how far they are located from the monitor. Hence, it does not account for the fact that pollution concentrations are likely to dissipate with distance.

The second interpolation approach addresses this issue by using a distance-weighted method. We generate a pollution surface for the entire study area using 100-meter-by-100-meter grid cells. We then assign to each grid cell a distance-weighted average of the readings from the three closest monitors. The neighborhood air quality measure is then computed by averaging the grid values within each PUMA. The two interpolation approaches lead to similar neighborhood ozone and PM-10 concentrations. We follow Sieg et al. (2004) and use the pollution levels from the closest monitor interpolation approach in the estimation of household preferences and the computation of welfare benefits.

4.3 Characterizing the Residential Location

We characterize the household's residential location choice alternatives in terms of 4037 discrete housing types. These are also referred as housing products. Each housing type is a collection of housing units that are located within the same neighborhood and have identical observed characteristics. Housing types are defined in terms of six variables: ownership status, number of bedrooms, dwelling type, built after 1980, built during 1960s or 70s and PUMA. The first five variables represent the housing characteristics for each housing type. We characterize the rental price of a given housing type *h* as the average of the rental prices for all units of type *h*. This is similar to the approach used by Berry et al. (1995) to obtain average prices of car products. The neighborhood characteristics for each housing type are given by the characteristics of the PUMA.

The ownership status is defined as either renter-occupied or owner-occupied. The number of bedrooms ranges from 0 to 5. The dwelling type is defined as either single-family or multi-family. The variables "built after 1980" and "built during 1960s or 70s" are binary variables that equal one if true and zero otherwise. Lastly, the study area contains 79 neighborhoods. These six categories provide a total of 7,584 ($2 * 6 * 2 * 2 * 2 * 79$) possible housing types. The actual number of combinations that exist in the study area is much smaller. We obtain a total of 4037 distinct housing products. This is because some of the 7,584 possible housing types do not exist in the data. For example, in a given neighborhood there are eight possible types of four-bedroom multi-family units. However, some neighborhoods contain no multi-family four-bedroom units. As a result these neighborhoods will have zero, instead of eight, types of four-bedroom multi-family units.

Using housing types rather than housing units to characterize residential locations significantly reduces the number of alternatives in the housing market while still providing a complete span of the product space. This has a direct implication for the identification of preference parameters in the first stage of the estimation. Indeed, a necessary requirement for the identification[17] of the first stage is that the number of observations be larger than the number of alternative-specific constants plus the number of interaction parameters. This requirement is not

[17] A discussion of identification issues is provided in the appendix.

met when housing units are used to characterize residential locations, as the number of observations (i.e., households) will equal the number of housing alternatives.

5 Estimation Strategy

We estimate the parameters of the household's indirect utility in equations (3.3) and (3.4). In section 5.1 we characterize the sampling framework used to generate the household sample and the choice set of sampled households. Section 5.2 discusses the details of the estimation strategy. Section 5.3 presents the results of the estimation.

5.1 Sampling Framework

Two issues arise in the empirical estimation of the household location choice model. The first issue regards how to draw the sample of households to be used in the estimation of the model. The sampling of households is necessary because it is not computationally feasible to estimate the model from the population of 171,000 households in the study area. The second issue pertains to the characterization of the relevant choice set for the sampled households. This is a classical issue in the estimation of discrete choice models (See, for example, McFadden, 1978 and Quigley, 1985).

5.1.1 Drawing the Household Sample

We devise a sampling scheme that allows using a smaller, yet representative sample of the households in the data. The sampling framework uses a stratified, choice-based sampling design.[18] In particular, we draw a 10 percent random sample of the households who choose each

[18] Ben-Akiva and Lerman (1985) provide a review of sampling theory and applications to the estimation of discrete choice models.

housing type. This produces the final sample of 17,894 households used to estimate the location choice model.

The choice-based sampling design does not produce a fully random sample of the household population. Indeed, it is easy to show that the average household characteristics from this sample will be biased estimates of the mean household characteristics in the population. An alternative to the choice-based sampling design would be to use a simple random sampling scheme. While a simple random sampling design produces independent observations, it does not guarantee that every housing type will be represented in the sample. This will likely be the case for housing alternatives that are chosen by very few households. In other words, the random sample may not produce households from those residential locations. In an attempt to provide a full characterization of the housing market, we opted to preserve the product space at the expense of the independence of household observations. We correct for the bias in the first-stage estimation, resulting from the choice-based sampling design, using the approach of Manski and McFadden (1981). This correction is explained below in the details of the estimation.

5.1.2 Determining the Choice Set of Sampled Households

The household's relevant choice set or feasible set of alternatives is an essential component of the estimation. A sampling approach is also used to construct the choice set. Potentially, one could set the household's choice set as the 4037 housing types in the sample. However, this would render the estimation computationally intractable. The reason is that the computational burden of the estimation grows linearly with the size of the household's choice set. An alternative is to construct the choice set by sampling a few alternatives from the full set of available alternatives. In particular, the household's choice set includes (i) the household's chosen residential location and (ii) a random sample of 20 residential locations from the remaining non-chosen alternatives. McFadden (1978) has shown that such a scheme will yield consistent parameter estimates for the multinomial logit model.

5.2 Estimation of Household Preference Parameters

The parameters (α, γ, β_0, β_1) of the household indirect utility function defined by equations (3.3) and (3.4) are estimated from a multinomial logit model. The estimation follows the two-stage approach proposed by Bayer et al. (2005). In the first stage we estimate (*H-1*) alternative-specific constants[19] (δ_h) and the household-specific taste parameters (α, γ, β_1) in equation (3.3). The second stage estimates the vector of mean taste parameters (β_0) using the estimated vector of alternative constants as the dependent variable in the regression specification given by equation (3.4).

5.2.1 Recovering the Household-Specific Taste Parameters (First Stage)

The alternative-specific constants (δ_h) and the household-specific taste parameters (α, γ, β_1) are obtained via maximum likelihood estimation (MLE). The indirect utility in equation (3.3) defines the household-specific multinomial choice probabilities given by:

$$P_{ih}(p, z_i, x; \delta, \alpha, \gamma, \beta_1) = \frac{\exp[\delta_h + \alpha \log(y_i - p_h) + \gamma d_{ih} + \sum_{kr} x_{hk} z_{ir} \beta_{1kr}]}{\sum_{m \in C_i} \exp[\delta_m + \alpha \log(y_i - p_m) + \gamma d_{im} + \sum_{kr} x_{mk} z_{ir} \beta_{1kr}]}, \qquad (5.1)$$

where C_i represents the choice set of household *i*. Given the household choice probabilities, the log-likelihood for the household choices observed in the data is defined as:

[19] Note: The H^{th} alternative constant is set to zero.

$$L(\delta,\alpha,\gamma,\beta_1)=\sum_i \sum_{h\in C_i} I_{ih} \log P_{ih}(p,z_i,x;\delta,\alpha,\gamma,\beta_1)\,, \qquad (5.2)$$

where I_{ih} is a dummy that equals 1 whenever household i chooses location h in the data. The estimates for the preference parameters (α, γ, β_1) and the choice-specific constants (δ) are then obtained via maximization of the log-likelihood L (δ, α, γ, β_1).

The closing conditions of the equilibrium model are implicitly enforced via maximization of the log-likelihood. As pointed out by Bayer et al. (2005), this can be observed from the first-order condition of the maximization problem. Differentiating the log-likelihood in (5.2) with respect to $\hat{\delta}_h$ yields:

$$\frac{\partial L}{\partial \hat{\delta}_h}=\sum_{i\in h}(1-\hat{P}_{ih})+\sum_{i\notin h}(-\hat{P}_{ih})=\sum_{i\in h}1-\sum_{i\in h}\hat{P}_{ih}-\sum_{i\notin h}\hat{P}_{ih}=s_h-\sum_i \hat{P}_{ih}=0\,, \qquad (5.3)$$

where $\hat{P}_{ih}$ is the estimated choice probability, s_h represents the sample housing supply for alternative h, and $i\in h$ indicates that household i chooses housing type h. Notice that equation (5.3) closely resembles the equilibrium condition in equation (3.9). It is indeed the sample equivalent of equation (3.9). Hence, the vector of alternative-specific constants which maximizes the log-likelihood also insures that the equilibrium condition in equation (3.9) holds for the sample.

The maximization of the log-likelihood in equation (5.2) with respect to the full set of parameters (δ, α, γ, β_1) is computationally demanding. This is because the dimension of δ (the vector of location-specific constants) is generally large. In this study, the housing market comprises a total of 4037 housing alternatives. This requires estimating 4036 alternative-specific constants in the first stage. As a result, maximizing the log-likelihood using standard search algorithms (i.e., Newton-Raphson, quasi-Newton or direct search) can be extremely slow and inefficient. A contraction mapping proposed by Berry et al. (1995) allows one to circumvent this computational burden by solving for the alternative-specific constants separately using the first-order condition in equation (5.3).

Equation (5.3) implicitly defines the vector of alternative-specific constants (δ) as a function of the household-specific taste parameters (α, γ, β_1) and the vector of housing-type supplies (s). This allows one to derive a concentrated version of log-likelihood as a function of (α, γ, β_1). The concentrated log-likelihood is given by:

$$L_c(\alpha,\gamma,\beta_1) = \sum_i \sum_{h \in C_i} I_{ih} \log P_{ih}(\delta(\alpha,\gamma,\beta_1),\alpha,\gamma,\beta_1)\,. \tag{5.4}$$

For given values of (α, γ, β_1) that maximize the concentrated log-likelihood L_c, we can obtain estimates of the alternative constants by solving the system in equation (5.3). The contraction mapping of Berry et al. (1995) provides a quick numerical solution to this system. It suggests solving iteratively for the location constants using the following recursive algorithm:

$$\delta_h^{t+1} = \delta_h^t - \log\left[\sum_i \hat{P}_{ih}(\delta_h^{\,t},\alpha,\gamma,\beta_1) \Big/ s_h\right]. \tag{5.5}$$

Berry et al. (1995) prove that the algorithm in equation (5.5) is a contraction mapping, which means that it is guaranteed to converge for any starting value of δ. Convergence generally occurs quickly. In our estimation, convergence of the contraction mapping usually occurs after 20 to 30 iterations. The computing time is between 5 and 10 seconds on Pentium 4 2Ghz PC stations.

The first stage estimation can be summarized as follows:

i. Set an initial guess for δ.

ii. Given δ, maximize the constrained log-likelihood in (6.4) with respect to (α, γ, β_1).

iii. Given the estimates of (α, γ, β_1), solve for δ using the contraction mapping in (6.5).

iv. Repeat (ii) and (iii) until the estimates converge.

It is easy to see that the above steps solve the system of first-order conditions for the unconstrained log-likelihood in equation (5.2). This implies that the estimates produced by this sequential estimation are indeed the MLE estimates of ($\delta, \alpha, \gamma, \beta_1$), which are unique given the global concavity of the multinomial logit log-likelihood.

5.2.2 Correcting for the Sampling Design

As discussed in the previous section, the choice-based sampling approach does not produce a random sample from the household population. As a result, additional steps need to be taken to ensure that the first-stage MLE estimates are consistent. It turns out that the log-likelihood in equation (5.2) represents a special case which requires only a minor correction to achieve consistency. In fact, it has been shown (McFadden and Manski, 1981) that the MLE estimates for (α, β_1) are consistent as long as (i) the choice model is a multinomial logit and (ii) the model contains a full set of alternative-specific constants (Ben-Akiva and Lerman, 1985). Both of these conditions are satisfied by the log-likelihood in equation (5.2). In addition, a minor correction will ensure the consistency of the alternative constants when the sampling design is such that each choice alternative is a stratum and the population share of each stratum is known. The consistent estimate of δ_h is obtained as:

$$\hat{\delta}_h = \hat{\delta}_h^{mle} - \ln\left({w_h}/{W_h} \right). \tag{5.6}$$

Where w_h is the fraction of the sample drawn from stratum h, and W_h represents the population share of stratum h. For the sampling design described in the previous section, each housing type represents a stratum. Therefore w_h is the ratio of the number of households drawn from housing type h to the total number of households in the sample. W_h is the proportion of the household population choosing each housing type h.

5.2.3 Estimating the Mean Taste Parameters (Second Stage)

In the second stage, the mean taste parameters (β_0) are estimated via ordinary least-squares (OLS). We regress the vector of alternative-specific constants estimated in the first stage on the housing and neighborhood attributes. The regression specification is given by:

$$\hat{\delta}_h = \sum_k x_{hk} \beta_{0k} + \xi_h . \tag{5.7}$$

The underlying assumption of the second-stage regression is that the housing and neighborhood attributes in x_h are uncorrelated with the unobserved attributes of the residential location. That is, they must be exogenous or at least determined prior to the revelation of the household's valuation for the unobserved attribute (Nevo, 2000). A potential endogeneity problem may be due to unobservable neighborhood attributes that may be correlated with neighborhood air quality. Bayer, Keohane and Timmins (2005) address this issue by constructing an instrument for neighborhood PM-10 air pollution that uses panel data. In particular, they compute the PM-10 measure, for a location *j*, using changes in PM-10 levels originating from sources outside location *j*. Though we recognize the potential endogeneity of the neighborhood ozone measure, the fact that we have a small number of neighborhoods (79) limits our ability to construct reliable instruments. However, robustness checks suggest that the endogeneity of the PUMA-level ozone measure is not a severe problem. We return to this issue below in the estimation results.

Differentiated product models (see e.g., Berry et al., 1995 and Bayer et al., 2005) have used an instrumental variable (IV) approach to deal with the potential endogeneity problem that arises when the housing price enters the second stage. This endogeneity is caused by the fact that housing prices are likely to be correlated with unobserved characteristics of residential locations. However, we do not instrument for housing prices as they do not enter the second-stage regression. Our model does not treat housing prices as attributes of residential locations. Rather, housing prices enter the first-stage estimation as part of the household's budget constraint. The first-stage maximum likelihood estimation does, however, assume that the household's expenditure on non-housing goods, i.e., the term (y-p), is uncorrelated with the household-specific random error term (ε_{ih}).

5.3 Estimation Results

We estimate the specification of the household's indirect utility function in equations (4.3) and (4.4). As explained in the data section, we use ozone concentrations to characterize air pollution in 1990. Due to the high correlation between the household characteristics we only estimate a limited set of interaction parameters in the first stage.

5.3.1 Parameter Estimates

Table 6 summarizes the results of the estimation.[20] Model 1 estimates the benchmark specification which is used in the welfare estimation. The other models provide robustness checks which are described below. The household-specific taste parameters estimated in the first stage are all significant. The interaction parameters also have the expected signs except for the interaction parameter between crime and household income. We find that households with higher income levels have a higher WTP for air quality, which is in accordance with the hypothesis that air quality is a normal good. We also find that larger households are willing to pay more for additional bedrooms. Households with college-educated heads tend to have a stronger preference for school quality. This is in accordance with the idea that more educated people place a higher value on the quality of their children's education. Households prefer residential locations that are within their employment zone, which is consistent with the notion that households dislike commuting.

The positive and significant interaction between the log of crime and household income is contrary to our intuition. We would tend to expect that public safety is a normal good. This means that households with a higher income would want to have more public safety and hence be willing to pay more. This would imply a negative sign for the interaction of crime with income. As described in section 4, the crime variable is quite noisy as crime rates are only available at the city level. Also, as Table 4 shows, there is not enough variation in the crime

[20] Except for the term Log(y-p), all the household-level interaction variables have been demeaned.

variable across neighborhoods. These factors may contribute to the counterintuitive interaction effect between crime and income.

The second-stage taste parameters also generally have the expected signs. On average, households are found to prefer more bedrooms, owner-occupied dwellings, single-family dwellings, better school quality and coastal communities. The second-stage ozone coefficient is not statistically different from zero at the 10 percent level. The mean taste for ozone can be obtained by multiplying the ozone-income interaction coefficient, -0.02, by the mean of Log(y-p) in our sample, 8.

The estimated ozone coefficient implies a mean marginal willingness to pay (MWTP) of $62 for a 1 percent reduction in the 1990 average ozone concentration. We follow Sieg et al. (2004) by reporting the MWTP for a 1 percent reduction in the 1990 ozone levels. This allows comparing the MWTP estimates with estimates from previous studies. Sieg et al. (2004) report a MWTP of $61 for a 1 percent reduction in the 1990 average ozone concentration.[21] We also find a significant variation in MWTP across households. For example, the MWTP for a 1 percent reduction in ozone for households in the highest income quartile (top 25 percent) is $130 compared to only $8 for households in the lowest-income quartile.

5.3.2 Robustness Checks[22]

Endogeneity of Neighborhood Air Pollution

As discussed in the previous section, the estimate of ozone pollution in the second-stage regression is likely to be endogenous as a neighborhood's ozone level may be correlated with unobserved neighborhood socioeconomic variables that enter the error term (ξ_h). As a result, the estimated mean taste parameter for air pollution may be biased and inconsistent. The direction of

[21] A detailed comparison of the MWTP estimates with the literature is provided in the appendix.
[22] Additional robustness checks are provided in the appendix.

this bias is to make the coefficient less negative, as air pollution will generally be positively correlated with neighborhood characteristics, such as share of low-income households and share of ethnic minorities, which are often disliked by high-income households. This could explain the positive estimate of ozone pollution in the second-stage regression.

As explained previously, the small number of neighborhoods in the data limits the construction of reliable instruments. However, we perform a simple robustness check for the endogeneity problem that would result from the correlation between neighborhood ozone level and unobserved neighborhood characteristics. This involves estimating the second-stage OLS regression without the proportion of Hispanics. The assumption is that the unobserved neighborhood socioeconomic variables are correlated with the proportion of Hispanics in the neighborhood. Hence, if the ozone level is correlated with unobserved socioeconomic characteristics, removing the neighborhood proportion of Hispanics from the second-stage regression should significantly lessen the bias in the estimated ozone mean taste parameter. Model 1a of Table 6 reports the results from the alternate regression specification. We find that the estimated ozone coefficient remains positive and insignificant. The magnitude of the coefficient is also roughly the same in Model 1 and Model 1a. We should again note that the mean taste for ozone remains negative, as the ozone-income interaction coefficient is the same across models 1 and 1a.

Robustness Checks with Respect to the Crime and Employment Variables

We mentioned previously that the crime variable is quite noisy as crime rates are only available at the city level. One may wonder whether the noisiness in the crime variable may significantly affect the estimates of the taste parameters for the other neighborhood variables. Model 2 of

Table 6 runs the estimation without the crime variable in both the first and second stages. The estimated parameters from this model are very similar to the estimates in Model 1.

The estimated taste parameter for the household's preference for locations that are within its employment zone is significantly large in absolute value compared to the other taste parameters. It is possible that the employment dummy may also be capturing households' preferences for other unobserved neighborhood characteristics. To the extent that this is the case, one may wonder if the presence of the employment dummy significantly distorts the estimated coefficient for ozone in both the first and second stages. As a robustness check, we run the estimation without the employment dummy in the first stage. The results are reported in Model 3 of Table 6. Except for the coefficients involving the crime variable, the remaining estimated parameters are similar to those in Model 1.

Alternative Characterization of Residential Locations

We explained in section 4.3 that the residential locations are characterized in terms of housing types, rather individual housing units. This not only considerably reduces the computational burden of the estimation, but also plays a key role in the identification and asymptotic properties of the estimates (see appendix). When residential locations are characterized in terms of individual housing units, the alternative constants may not be identified since $N < H + k - 1$.[23] One would essentially be trying to recover more parameters than the number of observations in the first-stage estimation.

Model 4 of Table 6 estimates the household preference parameters by characterizing residential locations using housing units. This is the approach used by Bayer et al. (2005). The sample of housing alternatives is formed by taking a random subsample of H (= 17,894) housing

[23] Here H is the number of housing alternatives and k represents the number of interaction parameters to be estimated in the first stage.

units from the 171,000 houses in the 1990 PUMS data for the study area. The household sample is given by the households choosing the H sampled housing units (i.e., $N = H$). This means that the first stage will involve estimating N-1 alternative-specific constants plus k interaction parameters from the location choices of N households. Hence, there are not enough observations to explain all the parameters in the first stage estimation. This is reflected by the likelihood ratio test result for the first stage estimation. The joint null hypothesis that the estimated alternative constants are all zero cannot be rejected.

6 The Benefits of the 1990 Clean Air Act Amendments

6.1 Simulation of the Counterfactual Locational Equilibrium

Induced price changes that result from the re-sorting of households are obtained by simulating the counterfactual equilibrium which would have emerged in 1990 if air quality levels were identical to those observed in 2000 while all other housing attributes and household characteristics remained at their 1990 levels. The counterfactual equilibrium is given by the new set of housing prices and the resulting household location choice probabilities which solve the market equilibrium condition in equation (3.9). Residential location demands are calibrated using the estimates of the preference parameters entering the household indirect utility function. The counterfactual equilibrium captures the changes in the air quality that occurred in the Los Angeles area between 1990 and 2000. Other factors characterizing the Los Angeles area housing market, such as population, household income and housing supply, are not allowed to change in this simulation. A detailed description of the simulation model is provided in the appendix at the end of the paper.

Figure 8 shows the PUMA-level average housing price changes in the counterfactual 2000 equilibrium. There are substantial housing-price changes across the study area, which suggests that the air quality changes that occurred between 1990 and 2000 led many households to change their location choices. We find that housing prices are lower in the counterfactual equilibrium in the areas with below average air quality improvements. These were also areas with the highest air quality levels in 1990 (see Figure 1). Average housing prices fell by as much as 13 percent in those areas. On the other hand housing prices in the counterfactual equilibrium are higher in the areas that experienced above average air quality improvements in 2000. These were areas with the highest ozone levels in 1990. Housing prices rose by as much 8 percent in those areas.

6.2 Welfare Measurements

We characterize and estimate Hicksian welfare measures which are derived from a random utility function with non-linear income effects. The household-level Hicksian welfare measure for the changes in air quality is obtained using the principle of compensating variation (*CV*). The compensating variation for an air quality change is defined as the reduction in the household's income which is such that the household's maximum utility after the change equals the maximum utility before the change. Hence, by definition, the compensating variation will be negative for an air quality improvement and positive for a reduction in air quality.

For the utility function (v_{ih}) defined by equation (3.1), the household compensating variation for the air quality improvements that occurred in the Los Angeles area is implicitly defined by

$$V_{ih}(y_i - p_h^0, x_{1h}^0, x_{2h}^0, \varepsilon_{ih}) = V_{ij}\,(y_i - p_j^1 - CV_i, x_{1j}^1, x_{2j}^0, \varepsilon_{ij})\,, \qquad (6.1)$$

where $V_{ih} = \underset{h}{Max}\{v_{ih}\}$. The superscript zero indicates the 1990 market conditions, and the superscript one indicates the market conditions after the air quality changes. For ease of exposition, the attribute vector is broken into two components. x_{1h} represents the air quality level at location h, and x_{2h} is a vector capturing all other attributes of the residential location.

The household level *CV* measure defined by equation (6.1) is a random variable as it is a function of the unobserved taste error ε. Hence the welfare measure that is of interest to policy analysis is the expected value of the household level compensating variation over the distribution of the unobserved taste error ε. We define this expectation as:

$$ECV = E[CV \mid (y, p^0, p^1, x^0, x^1, \varepsilon)] \qquad (6.2)$$

The expectation *ECV* will characterize households' average *WTP* for the air quality changes across the Los Angeles area.

A general closed form expression for *ECV* does not exist for the indirect utility function in equation (3.3). This is because in certain cases the *CV* measure may be a nonlinear function of the stochastic error term ε. As a result its expectation, which requires integrating out the nonlinear error term, cannot be characterized explicitly. McFadden (1999) suggested a general simulation approach for recovering the exact *ECV*. In this study, we adopt the simulation approach of McFadden to obtain the average and income distributional welfare impacts of the 1990 CAAA.

Direct vs. Equilibrium Welfare Measures

For the purpose of evaluating the benefits of the changes in air quality across the Los Angeles area two welfare measures are of interest. The first measure asks what households are willing to pay for the change in air quality at their residence, holding housing prices and all other attributes fixed. We will refer to this welfare measure as the direct WTP measure (*CV^d*) since it can be recovered directly from the indirect utility function. This measure is sometimes referred as the partial equilibrium welfare measure. For our random utility model, CV^d is implicitly defined by:

$$V_{ih}(y_i - p_h^0, x_{1h}^0, x_{2h}^0, \varepsilon_{ih}) = V_{ih}(y_i - p_h^0 - CV_i^d, x_{1h}^1, x_{2h}^0, \varepsilon_{ih}), \tag{6.3}$$

where the notation is similar to that used in equation (6.1).

The direct WTP measure does not, however, provide a complete picture of the welfare impact of the changes in air quality across the Los Angeles area. Bartik (1988) shows that CV^d provides a lower bound to the full, i.e. equilibrium, welfare impact of the air quality changes. We define the equilibrium welfare measure (CV^e) as the WTP measure which takes into account the induced changes in housing prices that occur as households change their residential location choice. It is given by:

$$V_{ih}(y_i - p_h^0, x_{1h}^0, x_{2h}^0, \varepsilon_{ih}) = V_{ij}(y_i - p_j^1 - CV_i^e, x_{1j}^1, x_{2j}^0, \varepsilon_{ij}). \tag{6.4}$$

The household's residential location choice *j* in the *ex-post* equilibrium differs from the location *h* in the benchmark equilibrium. This indicates that the household might change its residential location choice as a result of the change in air quality. CV^e is sometimes referred as the general equilibrium welfare measure.

6.3 Welfare Impacts of the 1990 CAAA

Our analysis of the benefits of the 1990 CAAA focuses on the changes in neighborhood ozone levels between 1990 and 2000. The neighborhoods of the Los Angeles area experienced significant reductions in ozone levels during the years that followed the 1990 CAAA. Table 7 summarizes the changes in ozone levels for the neighborhoods in our sample. The neighborhood average ozone concentration fell by nearly 21 percent between 1990 and 1995. By the year 2000, the average reduction in ozone levels was close to 40 percent. The changes in ozone levels also varied across the area. The neighborhoods of Los Angeles and San Bernardino counties experienced the greatest ozone reductions between 1990 and 2000, while Orange and Riverside counties had the smallest average fall in ozone levels.

The neighborhood ozone changes for our sample differ slightly from the changes in ozone levels used by Sieg et al. (2004). In Orange County, for instance, our neighborhood ozone reductions between 1990 and 1995 were 4 percent lower than the reductions observed by Sieg et al. The slight divergence in ozone changes can be attributed to the differences in neighborhood geography. This study characterizes neighborhoods with PUMA boundaries while Sieg et al. use school district boundaries to characterize neighborhoods.

6.3.1 Results

Mean Welfare Impacts

Table 8 presents the mean welfare impacts of the CAAA from 1990 to 2000. These are the exact welfare measures obtained via McFadden's simulation approach. The first row provides the overall results for the study area. The second group of rows provides the county-level results. The last two groups of rows provide results for selected neighborhoods. In the third set of rows neighborhoods are ranked by their average 1990 income level and the mean welfare results are presented for the 1st, 50th and 99th percentiles. In the last set of rows we rank neighborhoods by their 1990 ozone level and present the mean welfare results for the 1st, 50th and 99th percentiles.

The welfare results suggest that, on average, the 1990 CAAA provided significant benefits to the households of the Los Angeles metropolitan area. We estimate that the reductions in ozone levels between 1990 and 2000 provided an average equilibrium welfare benefit of $1,829 to the households of the Los Angeles area. This benefit represents 4 percent of the annual average household income in 1990. As conceptually predicted by Bartik (1988) and demonstrated by Sieg et al. (2004), direct welfare benefits, which do not account for induced changes in housing prices, underestimate the benefits of the air quality improvements. On average, equilibrium benefits were 32 percent higher than the direct benefit estimates.

The estimated mean welfare benefits vary somewhat across the counties in the sample. Average benefits are highest in Orange County and lowest in Los Angeles County. The mean equilibrium WTP for the ozone changes between 1990 and 2000 was $2,134 in Orange County. This compares with an average equilibrium benefit of $1,757 in Los Angeles County. The distribution of welfare benefits across counties tends to reflect equilibrium price effects across the counties. Orange County, which experienced a fall in housing prices, has a significantly larger average equilibrium WTP.

We find a significant variation in welfare gains across neighborhoods. The mean equilibrium benefit in the neighborhoods with the highest average income is nearly four times the mean equilibrium benefit in the poorest neighborhoods. This variation can be attributed to richer households that have a significantly higher MWTP for air quality compared to low-income households in our model. However, relative equilibrium gains are higher in the lower-income neighborhoods as evidenced by the ratio of equilibrium to direct benefits. Indeed, equilibrium benefits are 84 percent higher than direct benefits in the poorest neighborhoods, as compared to only 19 percent in the richest neighborhoods.

We also find that households originally located in the most polluted neighborhoods have, on average, lower equilibrium benefits than households originally located in the least-polluted neighborhoods. This variation can be attributed to the fact that the most polluted neighborhoods, which had above average ozone reductions, experienced an increase in housing prices. On the other hand, housing prices decreased in the least polluted neighborhoods as they generally had below average ozone reductions (an ozone increase in the case of the cleanest neighborhood).

Income Distributional Welfare Impacts

Table 9 presents the distribution of equilibrium welfare estimates across household income quartiles. The lowest income quartile is comprised of households with annual 1990 income below $20,000 dollars, whereas the highest income quartile includes households with annual income above $60,000. Income distributional benefits are provided for the study area as well as counties and neighborhoods.

Equilibrium benefits vary significantly across household income groups. Specifically we find that richer households generally have significantly higher benefits compare to households in the lower-income groups. This is true for the overall study area as well as within counties and neighborhoods. The variation in welfare gains across income groups is to be expected as the higher-income households have a significantly higher MWTP for air quality in our model.

We also find a somewhat significant variation in welfare gains across neighborhoods within each income group. For instance, high-income households who were located in neighborhoods with low and median air quality levels in 1990 have significantly higher benefits than the average high-income households. On the other hand, high-income households who resided in the dirtiest neighborhoods experience significantly lower benefits than the average high-income household in the study area. This disparity can be attributed to the fact that housing prices increased in the neighborhoods with the highest ozone levels in 1990 as a result of the above average air quality improvements in those neighborhoods.

Comparing Relative Welfare Gains across Income Groups

Figure 3 shows the mean WTP as a proportion of the household's income in 1990. The bar graphs characterize the distribution of relative welfare gains across income groups. The WTP estimates are obtained using McFadden's simulation approach. The distributional findings seem to differ between the direct and equilibrium welfare measures. While the direct welfare measure suggests that the richer households experienced higher relative welfare gains, the equilibrium welfare measure suggests that the distribution of relative benefits is fairly even across income groups. Hence ignoring equilibrium adjustments can also significantly alter the distribution of relative welfare gains.

The divergence of the distributional welfare picture in the direct and equilibrium approach can be explained from the relative difference between the two welfare measures which is also shown in Figure 3. This difference can be interpreted as the household's relative welfare gain from adjusting to a new location after the air quality changes. Figure 3 shows that the welfare gains from the equilibrium adjustments represent a larger share of income for low-income households. On the other hand, the direct welfare gains are larger for high-income households as they are willing to pay more for a marginal improvement in air quality. Hence the direct benefit measure will tend to misrepresent the distributional welfare impacts from large air quality changes.

6.3.2 Comparing with Previous Studies

To provide a comparison of our results with those of Sieg et al. (2004) we simulate the counterfactual equilibrium that would have resulted from the changes in ozone levels between 1990 and 1995. This is because, in their empirical analysis, Sieg et al. use the changes in ozone levels that occurred between 1990 and 1995. Table 10 reports the welfare results for the changes in ozone levels between 1990 and 1995. The results suggest that the reductions in ozone pollution between 1990 and 1995 provided an average equilibrium benefit of $896 to the households of the Los Angeles area. Similar to the welfare benefits from 1990 to 2000, there is a significant variation in the equilibrium benefits for 1995 across counties.

The last three columns of Table 10 report the overall and county-level mean benefit estimates from Sieg et al. (2004). The overall direct and equilibrium benefit estimates are substantially lower than the Sieg et al. estimates. The county-level benefit estimates also differ significantly. The county-level direct WTP estimates are consistently lower than the Sieg et al. estimates. The relationship between the equilibrium benefit estimates is, however, more complex. The equilibrium welfare estimates from this study are higher than the Sieg et al. benefit measures in Los Angeles and Orange counties. The relationship between the welfare measures is reversed in Riverside and San Bernardino counties. Sieg et al. also find that equilibrium adjustments in the 1995 counterfactual equilibrium resulted in average welfare losses for households in Riverside and San Bernardino counties. Our results, on the other hand, suggest that, on average, the equilibrium adjustments resulted in welfare gains for households in all four counties.

The disparity between our welfare estimates and those found by Sieg et al. can be due to a number of factors. First, the differences could emerge as a result of differences in the data. The fact that the two studies use a different characterization of neighborhoods (PUMA vs. school district) is likely to affect the welfare results. In addition, Sieg et al.'s average welfare benefit for the Los Angeles area includes Ventura County. We excluded Ventura County from our sample because the 1990 PUMA boundaries for that county were not mutually exclusive and hence did not meet the selection criteria (see section 4.2).

Second, the welfare results are likely to diverge from the Sieg et al. results because of the differences in the specification of households' location choices. The discrete choice characterization of households' location choices allows estimating household preferences that vary across income groups and educational levels. The preference estimates suggest that high-income households have stronger preferences for air quality relative to the average population. We also find that the average household population has a lower preference for school quality compared to college educated households. This contrasts with the Sieg et al. framework in which households are restricted to have the same preference ordering of neighborhoods with respect to neighborhood amenities. This is due to the fact that the marginal rate of substitution between community amenities is independent of the household's income and taste. In addition, the preference specification in this study naturally captures the geography of the housing market by allowing household preferences for locations to depend on the proximity to their employment location. We find that households have stronger preferences for housing alternatives that are located within their employment zone.

6.3.3 Limitations

We now discuss some limitations of the equilibrium welfare measures developed. The equilibrium welfare estimates in this study are based on the simulation of a counterfactual equilibrium which only accounts for air quality changes and induced housing price changes that result from the resorting of households. The actual welfare impacts of the 1990 CAAA should also account for changes in the housing supply, household income, household population and other salient changes which occurred between 1990 and 2000. These changes will likely affect the welfare benefits of the 1990 CAAA.

Using the higher household income levels in 2000 would likely lead to higher benefit estimates as high-income households have a higher marginal WTP for air quality. If the supply of housing is elastic with respect to price, accounting for housing supply adjustments would likely increase equilibrium benefits as the influx of new housing units would provide more choices to households. An increase in population is not likely to affect equilibrium welfare gains to the extent that the increased demand for housing result in higher prices everywhere.

The estimated equilibrium welfare measures could also be sensitive to the geographic definition of the housing market. We assume in this work that the Los Angeles area housing market comprises four counties: Los Angeles County, Orange County, Riverside County and San Bernardino County. One could argue, as in Sieg et al. (2004), that the Los Angeles area housing market also includes Ventura County. All else equal, a larger geographic area is likely to lead to higher welfare benefits as it would provide more choices to households.

The equilibrium welfare measures could also be sensitive to the geographic characterization of neighborhoods. This study uses the 1990 Census Public Use Microdata Areas (PUMA) to characterize neighborhoods. On the other hand, Sieg et al. (2004) use the 1990 school district boundaries to define neighborhoods. One could also characterize neighborhoods using smaller geographic units such Census tracts, Census blocks groups or Census blocks. Altering the geographic definition of neighborhoods is less likely to significantly affect the air quality measures as they generally do not vary much across small areas. As a result, welfare impacts of air quality changes are likely to be less sensitive to the characterization of neighborhoods.

The random utility specification in equation (3.3) also assumes away endogenous social interaction effects. Our utility function incorporates an exogenous social interaction effect. The social interaction effect is a result of households' homogeneous tastes for the proportion of Hispanics in the neighborhood. Incorporating endogenous social interactions in the household's utility could affect the equilibrium welfare estimates. For example, low-income renters could suffer welfare losses as increases in housing prices in their original neighborhoods force them to relocate to neighborhoods with less desirable attributes. An avenue for future research would be to explore empirically the extent to which the overall and distributional impacts of the 1990 CAAA are affected when endogenous social interactions are incorporated in the household's random utility function.

7 Conclusions

This study has developed a discrete choice equilibrium model to evaluate the benefits of the air quality improvements that occurred in the Los Angeles area between 1990 and 2000 as a result of the implementation of the 1990 Clean Air Act Amendments. The study has two main objectives. The first is to apply the discrete choice equilibrium framework (Anas, 1980, 1982) to the valuation of large environmental changes. The second objective is to evaluate the distributional welfare impacts of the 1990 CAAA in the Los Angeles area.

The empirical analysis suggests that the reductions in ozone concentrations across Los Angeles, Orange, Riverside and San Bernardino counties, provided an average equilibrium benefit of $1,800 to households. In contrast, average benefits are $1,400 when equilibrium price effects are not accounted, demonstrating that ignoring equilibrium effects will likely underestimate the benefits of large environmental changes. We find that the equilibrium welfare impacts of the 1990 CAAA in the Los Angeles area varied significantly across income groups.

Households in the highest income quartile experienced equilibrium benefits of approximately \$3,600 as compared to only \$400 for households in the lowest income quartile. The study also finds that ignoring equilibrium adjustments in housing prices can significantly alter the distribution of relative welfare gains (i.e., welfare gains as a proportion of household income). Indeed, welfare impacts that do not account for equilibrium effects suggest that high-income households have larger relative welfare gains compared to low-income households. However, when accounting for equilibrium adjustments, the distribution of relative welfare gains from the 1990 CAAA is fairly even across income groups.

References

Anas, A. 1980. “A Probabilistic Approach to the Structure of Rental Markets.” *Journal of Urban Economics* 7: 225-47.

Anas, A. 1982. *Residential Location Markets and Urban Transportation: Economic Theory, Econometrics and Policy Analysis with Discrete Choice Models*. Academic Press, New York.

Athey, S., and G. Imbens. 2007. “Discrete Choice Models with Multiple Unobserved Product Characteristics.” *International Economic Review* (forthcoming).

Bartik, T. 1988. “Measuring the Benefits of Land Improvements in Hedonic Price Models.” *Land Economics* 64(2): 172-83.

Bayer, P., N. Keohane and C. Timmins. 2005. “Migration and Hedonic Valuation: The Case of Air Quality.” NBER Working Paper: 12106.

Bayer, P., R. McMillan and K. Rueben. 2005. “An Equilibrium Model of Sorting in an Urban Market.” NBER Working Paper: 10865.

Bayer, P. and C. Timmins. 2005. “Estimating Equilibrium Models of Sorting Across Locations.” *Economic Journal* (Forthcoming)

Berry, S., J. Levinsohn and A. Pakes. 1995. “Automobile Prices in Market Equilibrium.” *Econometrica* 63(4): 841-90.

Ben-Akiva, M., and S. R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge.

Epple, D. and H. Sieg. 1999. “Estimating Equilibrium Models of Local Jurisdictions.” *Journal of Political Economy* 107(4): 645-81.

Manski, C. F. and D. McFadden. 1981. “Alternative Estimators and Sample Designs for Discrete Choice Analysis,” in C. F. Manski and D. McFadden, eds., *Structural Analysis of Discrete Data with Economic Applications*. Cambridge, Mass: MIT Press.

McFadden, D. 1973. “Conditional Logit Analysis of Qualitative Choice Behavior,” in: P. Zarembka, ed., *Frontiers of Econometrics*. New York: Academic Press.

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

___________. 1999. “Computing Willingness-to-Pay in Random Utility Models.” in J. Moore, R. Riezman, and J. Melvin, eds., *Trade, Theory and Econometrics: Essays in Honor of John S. Chipman*, London: Routledge.

Morey, E, R. D. Rowe and M. Watson. 1993. “A Repeated Nested-Logit Model of Atlantic Salmon Fishing.” *The American Journal of Agricultural Economics* 75(3): 578-92.

Nevo, A. 2000. “A Practitioner’s Guide to Estimation of Random-Coefficients Logit Models of Demand.” *Journal of Economics and Management Strategy* 9(4):513-548.

Palmquist, R. B. 2006. “Property Value Models,” in: Karl-Göran M. et al., eds., *Handbook of Environmental Economics, Volume 2.* Amsterdam: Elsevier North-Holland.

Palmquist, R.B. (1988). “Welfare measurement for environmental improvements using the hedonic model: the case of nonparametric marginal prices.” *Journal of Environmental Economics and Management* 15(3): 297–312.

Quigley, J. M. 1985. “Consumer Choice of Dwelling, Neighborhood and Public Services.” *Regional Science and Urban Economics* 15(1):41-63.

Rosen, S. 1974. “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition.” *Journal of Political Economy* 82(1): 34-55.

Sieg, H., V. K. Smith, H. S. Banzhaf and R. Walsh. 2004. “Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods.” *International Economic Review* 45(4): 1047-77.

Shadbagian, R., W. Gray and C. Morgan. 2004. “The 1990 Clean Air Act Amendments: Who Got Cleaner Air -- and Who Paid For It?” Mimeo, Clark University, Worcester, MA.

Smith, V. K. and J. Huang. 1995. “Can Markets Value Air Quality? A Meta-Analysis of Hedonic Property Value Models.” *Journal of Political Economy* 103(1):209-27.

Smith, V. K., H. Sieg, H. S. Banzhaf and R. Walsh. 2004. “General Equilibrium Benefits for Environmental Improvements: Projected Ozone Reductions under EPA’s Prospective Analysis for the Los Angeles Air Basin.” *Journal of Environmental Economics and Management* 47(3):559-84.

Takeuchi, A., M. Cropper and A. Bento. 2005. “The Welfare Effects of Slums Improvements: The Case of Mumbai.” Mimeo, University of Maryland, College Park, MD.

Timmins, C. 2007. “If You Can’t Take the Heat, Get Out of the Cerrado… Recovering the Equilibrium Amenity Cost of NonMarginal Climate Change in Brazil.” *Journal of Regional Science* 47(1):1-25.

U.S. Environmental Protection Agency. 2006a. "Clean Air Act: Full Text of 1990 Amendments." http://www.epa.gov/oar/caa/caaa.txt (accessed October 25, 2006).

U.S. Environmental Protection Agency. 2006b. "National Ambient Air Quality Standards." http://epa.gov/air/criteria.html (accessed October 31, 2006).

U.S. Environmental Protection Agency. 2006c. "8-hour Ground-level Ozone Designations – Region 9: State Designations." http://www.epa.gov/ozonedesignations/regions/region9desig.htm (accessed October 21, 2006).

U.S. Environmental Protection Agency. 1999. The Benefits and Costs of the Clean Air Act 1990 to 2010. *Report to Congress*.

Walsh, R. 2003. "Endogenous Open Space Amenities in Locational Equilibrium." Working Paper: 04-03. Department of Economics, University of Colorado, Boulder, CO.

Table 1: Average† Monitor Readings for Ozone and PM-10 [24]

		Study Area	Los Angeles County	Orange County	Riverside County	San Bernardino County
Ozone*	1990	0.144	0.150	0.116	0.137	0.154
	2000	0.097	0.089	0.078	0.111	0.109
Ozone Exceedances**	1990	36	37	11	33	47
	2000	3	2	0	5	6
PM-10 ***	1990	55.4	51.5	42.3	61.1	59.5
	2000	44.1	41.6	34.5	44.8	52.2

* Average top 30 1-hour daily maximum readings at a monitor during a year (parts per million).
** Number of days with a recorded violation the one-hour national standard for ozone.
*** Annual geometric mean (ug/m^3).
† The yearly reading for each monitor is obtained by computing a three-year centered average. For instance, the 1990 reading for monitor x is computed by averaging the readings for 1989, 1990 and 1991 at monitor x.

Table 2: Mean Household and Housing Characteristics in the 1990 PUMS

	Study Area	Los Angeles County	Orange County	Riverside County	San Bernardino County
Number of observations	170,955	119,726	28,209	5,642	17,378
Housing characteristics					
Monthly housing price ($)	749	709	956	725	707
1 if unit owned	0.51	0.47	0.58	0.63	0.63
Bedrooms	2.25	2.09	2.58	2.71	2.66
1 if built in 80s or 90s	0.19	0.15	0.24	0.43	0.32
1 if built in 60s or 70s	0.37	0.33	0.56	0.33	0.39
1 if single-family dwelling	0.62	0.58	0.66	0.77	0.76
1 if unit is within householder's employment zone	0.505	0.529	0.444	0.447	0.466
Household characteristics					
Monthly income ($)	4,098	3,943	4,945	3,860	3,926
1 if Asian and non-Hispanic	0.082	0.089	0.075	0.041	0.055
1 if Black and non-Hispanic	0.091	0.111	0.015	0.072	0.080
1 if Hispanic	0.237	0.262	0.147	0.189	0.224
1 if White and non-Hispanic	0.585	0.533	0.758	0.689	0.633
1 if children under 18	0.417	0.405	0.396	0.505	0.502
1 if married and has children under 18	0.015	0.014	0.015	0.014	0.017
1 if householder is 65 or older	0.16	0.17	0.13	0.12	0.13
1 if householder has college degree	0.35	0.33	0.44	0.29	0.32
Household size	2.99	2.97	2.95	3.14	3.16

[24] Source: California Ambient Air Quality Data. 2004 Data CD

Table 3: Mean Neighborhood (PUMA) Characteristics in 1990

	Study Area	Los Angeles County	Orange County	Riverside County	San Bernardino County
Number of observations	79	55	11	3	10
8th grade math score†	34.0	31.6	45.1	34.3	34.8
Crime (FBI index)	786.5	843.3	604.2	831.6	661.2
Elevation (meters)	200.7	172.9	63.2	345.6	461.8
PUMA is on Pacific coastline	0.114	0.091	0.364	-	-
Housing density (sq. km)	1,061.7	1,116.2	1,056.9	2,022.2	479.4
Ozone‡ (ppm)	0.146	0.143	0.109	0.177	0.198
Ozone Exceedances	32.94	29.58	12.11	51.46	68.80
PM-10 annual average ($\mu g/m^3$)	55.51	51.87	60.45	68.72	66.12

† School district average for 1994 CLAS. Math test scores have been normalized so they fall between 0 and 100.
‡ Annual average of top 30 daily 1hr maximum readings. PUMA is assigned the three-year centered average from the closest monitor.

Table 4: Within and between variation for Selected PUMA Characteristics in 1990

	Mean of PUMA Values	Std. of PUMA Means	Mean of within PUMA Std.
8th grade math score	34.0	35.5	5.7
Crime (FBI index)	786.5	770.1	631.9
Ozone†	0.146	0.040	-
Ozone‡	0.148	0.031	0.006
PM-10†	55.5	11.0	-
PM-10‡	53.2	7.4	1.4

† Interpolation method: PUMA is assigned closest monitor reading.
‡ Interpolation method: PUMA is assigned distance-weighted average of readings from three closest monitors.

Table 5: Correlation between Primary and Secondary Pollutants in 1990

	Ozone	PM-10	Nitrogen Oxide (NOx)	Sulfur Dioxide (SO_2)
Ozone	-	0.44	0.47*	-0.56**
PM-10	0.44	-	0.52*	-0.54

Note: * Significant at 5 percent level. ** Significant at 1 percent level.

Table 6: Estimation Results

	Model 1‡	Model 1a	Model 2	Model 3	Model 4
First Stage					
Log(y-p)	1.475**	-	1.499**	1.649**	2.052**
Ozone * Log(y-p)	-0.019**	-	-0.020**	-0.028**	0.01**
Bedrooms * Household size	0.066**	-	0.066**	0.066**	0.064**
Single family * Children under 18	0.227**	-	0.227**	0.227**	0.165**
Math * College educated head	0.309**	-	0.31**	0.244**	0.337**
Log crime * Log(y-p)	0.004**	-	-	-0.013**	0.026**
Within household's employment zone	1.989**	-	1.989**	-	2.194**
Log-Likelihood	-37,072	-	-37,072	-40,719	-47,733
Likelihood Ratio statistic (H_0: δ =0)	25,996	-	26,009	26,857	5,541
Likelihood Ratio p-value (H_0: δ =0)	0.000	-	0.000	0.000	0.999
McFadden pseudo-R^2	0.319	-	0.319	0.252	0.124
Observations	17,894	-	17,894	17,894	17,894
Second Stage OLS †					
Bedrooms	0.04*	0.041*	0.04*	0.044*	0.155**
Built after 1980	-0.594**	-0.594**	-0.594**	-0.596**	0.267**
Built in 60s or 70s	-0.172*	-0.171*	-0.173**	-0.169**	0.131**
Single-family dwelling	0.352**	0.346**	0.353**	0.349**	0.185**
Owned	0.054	0.057	0.053	0.04	0.044**
Math test score	0.139**	0.172**	0.153**	0.086*	0.092**
Log FBI crime index	0.0005	-0.0005	-	0.003**	-0.044**
Log elevation	0.016	0.035	0.007	-0.018	0.066**
PUMA is on Pacific coastline	0.342**	0.378**	0.327**	0.315**	0.167**
Log density	0.079	0.075	0.068	0.001	0.188**
Prop. of population Hispanic	-0.380*	-	-0.32*	-0.498**	-0.611**
Ozone	0.161	0.120	0.17	0.211	-0.095**
R^2	0.054	0.053	0.054	0.052	0.302
Observations	4,037	4,037	4,037	4,037	17,894

Notes:
** Significant at 1% level. * Significant at 5% level. † Standard errors are computed using White's robust covariance matrix.
‡ Model 1 : Benchmark specification used in the simulation and welfare analysis.
Model 1a: Estimates the second stage **without** the variable "proportion of Hispanics." This is intended to check the endogeneity of neighborhood ozone.
Model 2: Estimates the first and second stage **without** the "crime" variable.
Model 3: Estimates the first stage **without** the "employment" variable.
Model 4: Characterizes residential locations using individual houses instead of discrete housing types.

Table 7: Changes in Neighborhood Ozone Levels across the Los Angeles Area

	1990	1995	2000	% Δ 1990-95	% Δ 1990-95 (Sieg et al.)	% Δ 1990-2000
Study area	0.146	0.116	0.089	-20.8	-19.3	-38.9
Los Angeles County	0.143	0.110	0.086	-22.6	-20.8	-39.8
Orange County	0.109	0.094	0.076	-13.8	-18	-29.8
Riverside County	0.177	0.140	0.115	-20.6	-20.7	-35.2
San Bernardino County	0.198	0.162	0.115	-18.1	-16.3	-41.9

Table 8: Mean Direct (D) and Equilibrium (E) WTP* for the CAAA (1990-2000)

	Avg. 1990 Income	1990 Ozone	% Δ Ozone	1990 Avg. Price	% Δ Price	WTP_D	WTP_E	$WTP_{E/D}$
Study area (mean)	49,197	0.146	-36.1	748	0.14	1,386	1,829	1.32
Counties								
Los Angeles County	47,152	0.143	-37.6	728	0.17	1,325	1,757	1.33
Orange County	60,924	0.109	-23.8	926	-4.10	1,659	2,134	1.29
Riverside County	47,374	0.177	-34.4	687	1.02	1,299	1,764	1.36
San Bernardino County	48,096	0.198	-41.9	682	4.35	1,384	1,836	1.33
Neighborhoods by income levels								
1st percentile (lowest)	24,657	0.103	-46.8	455	-1.14	382	704	1.84
50th percentile	47,331	0.119	-40.4	805	-1.33	1,157	1,665	1.44
99th percentile (highest)	92,708	0.148	-48.7	982	2.57	2,378	2,837	1.19
Neighborhoods by ozone levels								
1st percentile (lowest)	65,135	0.058	30.0	1,000	-12.93	2,018	2,434	1.21
50th percentile	54,568	0.148	-43.7	822	1.41	1,462	1,832	1.25
99th percentile (highest)	39,979	0.212	-43.9	580	5.22	1,109	1,492	1.35

* Note: WTP is computed as the expected compensating variation (*ECV*). All WTP estimates computed using McFadden's simulation approach. WTP estimates are in annual 1990 dollars.

Table 9: Distribution of Equilibrium WTP* for the CAAA (1990-2000)

	Avg. 1990 Income	% Δ Ozone	% Δ Price	WTP Income < 20k	WTP Income 20k - 37k	WTP Income 38k - 60k	WTP Income > 60k
Study area (mean)	49,197	-36.1	0.14	441	1,019	1,706	3,634
Counties							
Los Angeles County	47,152	-37.6	0.17	433	1,009	1,682	3,638
Orange County	60,924	-23.8	-4.10	518	1,058	1,707	3,774
Riverside County	47,374	-34.4	1.02	384	1,053	1,796	3,133
San Bernardino County	48,096	-41.9	4.35	433	1,017	1,812	3,510
Neighborhoods by income levels							
1st percentile (lowest)	24,657	-46.8	-1.14	409	850	1,439	2,325
50th percentile	47,331	-40.4	-1.33	392	1,075	1,695	2,566
99th percentile (highest)	92,708	-48.7	2.57	479	909	1,590	4,505
Neighborhoods by ozone levels							
1st percentile (lowest)	65,135	30.0	-12.93	577	1,090	1,759	4,015
50th percentile	54,568	-43.7	1.41	388	964	1,790	4,341
99th percentile (highest)	39,979	-43.9	5.22	541	845	1,527	2,761

* Note: WTP is computed as the *ECV*. All WTP estimates computed using McFadden's simulation approach. WTP estimates are in annual 1990 dollars.

Table 10: Direct and Equilibrium WTP for the CAAA (1990-1995)

	Discrete Choice Equilibrium Approach			Epple-Sieg Equilibrium Approach (Sieg et al., 2004)		
	WTP_D	WTP_E	$WTP_{E/D}$	WTP_D	WTP_E	$WTP_{E/D}$
Study area	589	896	1.52	1,210	1,371	1.13
Los Angeles County	568	866	1.52	1,472	1,556	1.06
Orange County	698	1,029	1.47	901	1,391	1.54
Riverside County	526	858	1.63	834	372	0.45
San Bernardino County	576	891	1.55	738	367	0.50

Figure 1: 1990 and 2000 Ozone[25] Concentrations for the Greater Los Angeles Area

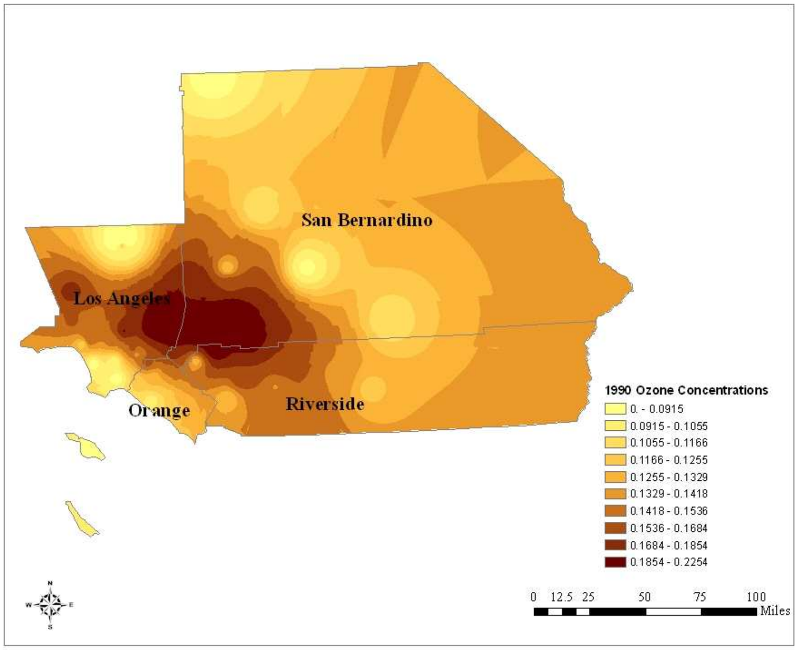

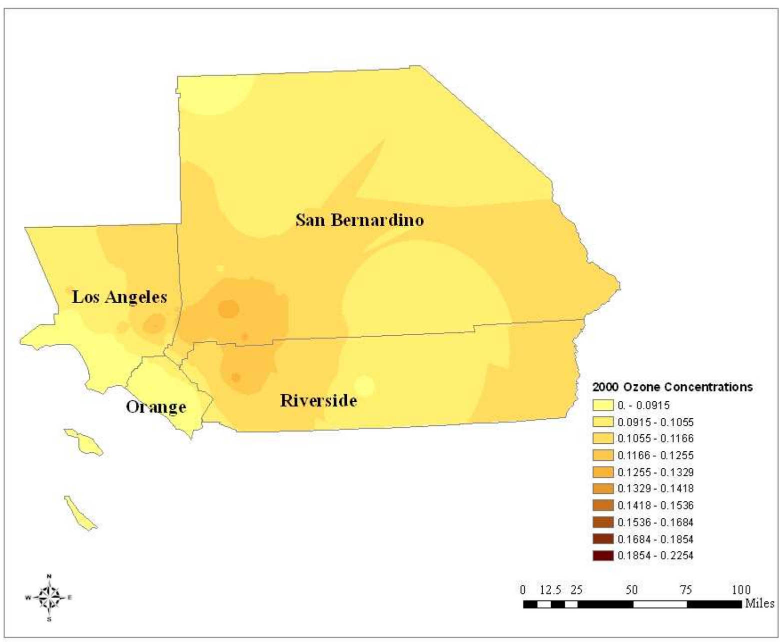

[25] Ozone concentrations are obtained via interpolation. We generate a pollution surface for the entire study area using 100-meter-by-100-meter grid cells. We then assign to each grid cell a distance-weighted average of the readings from the three closest monitors.

Figure 2: Percent Housing Price Changes in Counterfactual Simulation (PUMA average)

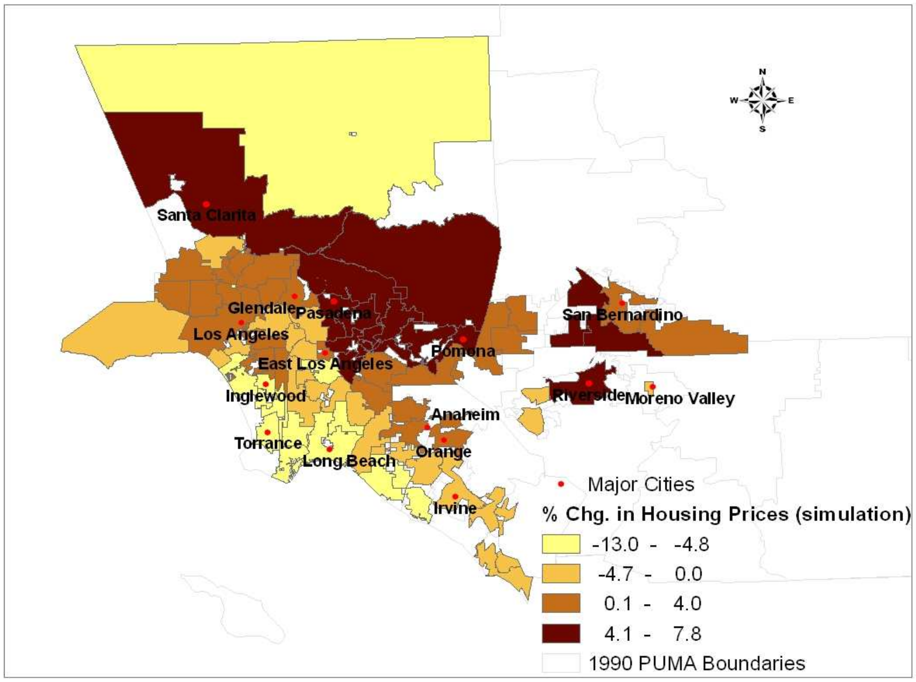

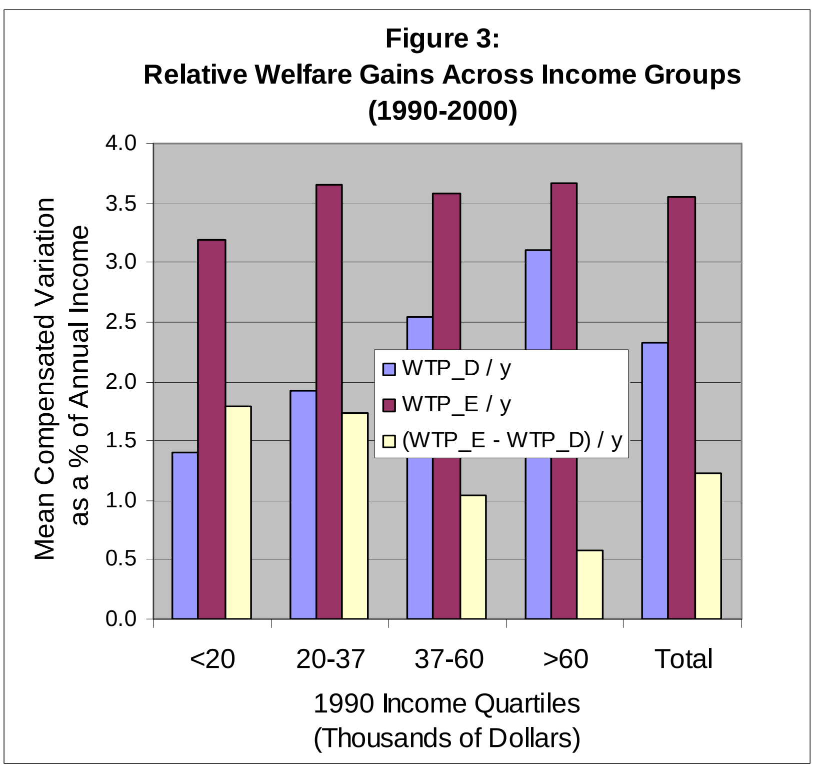

Appendix for "Equilibrium Welfare Impacts of the 1990 Clean Air Act Amendments in the Los Angeles Area" by Constant Tra

Section A of this appendix describes the procedure used to arrive at a single measure of housing price for owned and rental housing units. Section B describes the remainder of the neighborhood attribute data. Section C discusses the asymptotic properties of the estimated household preference parameters. Additional robustness checks for the estimated parameters are provided in section D. Section E compares the marginal willingness to pay (MWTP) values, implied by the parameter estimates, with other studies in the literature. The final section describes the simulation model.

A. Computing the rental price of housing across tenure

The housing price is a key characteristic which determines the sorting of households in our model. In the Census data the price of a house is reported as the owner's assessment of the market value, in the case of an owner-occupied unit, or the monthly rent in the case of a renter-occupied unit. To arrive at one price variable which will characterize both owner- and renter-occupied units we follow the approach of Bayer et al. (2005) by converting the market value of owner-occupied units to a monthly rental rate. Before describing this procedure we address some potential issues with the reported market value and monthly rent.

Value of Owner-Occupied Housing

A number of issues must be addressed when using the house value reported in the Census long form. The first issue relates to the fact that the housing price reported in the Census long form is based on the owner's own assessment of the market value. This assessment may not always reflect the true market value of the house, as most owners may either report the price of the house at the time of purchase or simply misrepresent the true market value of the house. The second issue regards the fact that the housing values reported in the 1990 Census are top-coded at $500,000. Because housing prices in California are generally higher than the remainder of the United States, we would expect to see a higher occurrence of binding top-codes. According to the 2000 Census 11.4 percent of houses in California where reported at a value of $500,000 or more compared to only 2 percent for the overall United States. In our 1990 sample approximately 8 percent of the houses have top-coded values.

To address these issues, we construct a predicted value for each house by making use of the property tax payment reported for each owner-occupied housing unit. The predicted value makes use of the California law (Proposition 13) that requires the property tax to equal either 1 percent of the transaction price of the house at the time the current owner bought the property or the value of the house in 1978. The predicted market value of each owner-occupied house is obtained by regressing the log of the reported house value on the estimated transaction price, i.e. 100 times the property tax, and a set of dummy variables for the year that the house was purchased. The regression specification is given by:

$$\log(p_h) = \alpha_1 \log(T_h) + \alpha_2 y_h + \varepsilon_h \,. \tag{1}$$

Where p_h represents the reported market value, T_h represents the estimated transaction price and y_h is a set of year dummies.

If the reported values were true, and all houses were identical except for the year of sale, then α_1 would equal 1 and α_2 would represent how much the house has appreciated in value. If, on the other hand, long-time owners tend to underreport the value of their house then α_2 would underrepresent the appreciation of the house in the market. In this case, the predicted value of the house from equation (1) should be a conservative estimate of the true market value. We replace the reported value for each house with our computed estimate whenever the latter exceeds the former, which would represent a case of significant underreporting on the part of the owner. In the actual implementation we allow the parameters to vary across subregions of our study area by running the regression in (1) for each of the three metropolitan statistical areas (MSA) in the study area. These are, Los Angeles-Long Beach, Orange County and Riverside-San Bernardino.

To correct for the bias in the house values, resulting from top coding, we use the following procedure. First, we estimate equation (1) using only the sample of houses whose values do not equal the top-code. We then use the estimated parameters to predict the market value for the houses with reported top-coded values. The estimated regression specification is reported in Table A.1.

Reported Housing Rents

As in the case of reported owner-occupied house values one may expect that reported monthly rents of renter-occupied units may not represent a fair assessment of the true market rent. This is likely to be true when the resident has lived in the house for a long period of time. In this case, we may expect that the reported rent will be an understatement of the true market rent. This could be either a result of rent controls or implicit tenure discounts. To correct this issue we compute an adjusted market rent by regressing the log of the reported market rent on a set of dummies characterizing the tenure of the current owner as well as a vector of housing characteristics. The regression specification is given by:

$$\log(p_h) = \beta_1 y_h + \beta_2 X_h + \omega_h . \qquad (2)$$

Where y_h is a dummy variable representing the year the current renter moved into the unit, and X_h is a set of housing and neighborhood characteristics for the house. As in the case of housing values we run this specification for each of the three MSAs in our sample. The estimated regression specification is reported in Table A.2. The parameter β_1 in equation (2) represents the tenure discount in a given PUMA. The corrected rent is then obtained as:

$$p_h^{corrected} = \exp[\log(p_h) - \beta_1 y_h].$$

Imputing the Rental Value of Housing across Units

In order to arrive at a comparable measure of housing price for both owner- and renter-occupied units, we convert owner-occupied house values into monthly rents using the approach described in Bayer et al. (2005). Poterba (1992) provides the theoretical foundation for this approach. Sieg et al. (2004) also use this approach to develop a price index for each housing unit in their sample. To convert housing values into monthly rents, we regress the log of the housing price (house value or monthly rent) on a dummy variable (O_h), indicating whether the unit is owner occupied, and a set of structural housing characteristics (X_h).

$$\log(p_h) = \gamma_1 O_h + \gamma_2 X_h + \upsilon_h \qquad (3)$$

We run this specification for each of the three MSAs (Los Angeles-Long Beach, Orange County and Riverside-San Bernardino) in our sample. The estimated regression specification is reported in Table A.3. The parameter γ_1 represents the ratio of house values to rents for each MSA, controlling for structural characteristics of housing units. This is the user-cost of owner-occupied housing as defined by Poterba (1992). We use this ratio to convert owner-occupied house values to a corresponding monthly rent.

To summarize, there are three sets of adjustments that are used to characterize the price of housing across owner-occupied and renter-occupied units. The first adjustment accounts for the fact that the house values contained in the Census data are self-reported and top coded. The second adjustment addresses the fact that housing rents contained in the Census data may misrepresent the true market rent. The final adjustment deals with converting owner-occupied housing values into monthly rents.

B. Other Neighborhood Data

In addition to air quality, we collect data on other neighborhood amenities that households may value. These include school quality, crime and racial composition. The racial composition of the PUMA is characterized by the proportion of Hispanics. Finally, three variables are used to control for unobserved factors that may affect the level of air pollution in a neighborhood. These are mean elevation of the neighborhood, the proximity of the neighborhood to the Pacific coastline and the housing density of the neighborhood.

School Quality

Because California state law limits expenditures of local school districts, a more reliable measure of school quality would be one that is based on academic performance outcomes rather than expenditures (Sieg et al. 2004). The California Department of Education (CDE) administers standardized tests that are used to monitor the academic performance of public schools. In the early 1990s the California Learning Assessment System (CLAS) was administered to public schools throughout the State of California. The 1994 CLAS provides a measure of students' academic performance in math, reading and writing. More recent academic performance test scores are the Academic Performance Index (API) and the STAR report.

We use the school district average 8th grade math score from the 1994 CLAS as the measure of school quality in 1990. Ideally one would want to use the 1989 CLAS data. Unfortunately this dataset is no longer available. The neighborhood school quality variable is computed by using a weighted average of the scores for all the school districts that intersect the PUMA. We use the area of the school district which intersects the PUMA as weight. For instance, suppose PUMA j has total area A and overlaps area $a(x)$ of school district x and area $a(y)$ of school district y. Then the school quality level for PUMA j is computed as:

$a(x) \cdot score(x)/A + a(y) \cdot score(y)/A.$

Figure B.1 provides a map of the neighborhood-level school quality data. The large cluster of neighborhoods with the worst school quality levels is part of the Los Angeles unified school district (LAUSD). The LAUSD is one of the largest school districts in the United States and the largest in the State of California.

Crime Rate

Currently, the most disaggregated crime data for California are provided by the Criminal Justice Statistics Center (CJSC) from the Office of the California Attorney General. The CJSC compiles statewide, county and city crime statistics and publishes them every year in the *Criminal Justice Profiles*. The crime variable used in this study is the FBI crime index for each jurisdiction in 1990. The FBI crime index reports the number of crime occurrences per 10,000 populations. The neighborhood crime rate is computed using the same weighting average method used to compute the neighborhood school quality. The crime data are not as reliable as the school quality data because they are only provided at the jurisdiction level and not all of the study area is incorporated. A map of the neighborhood crime levels in 1990 is shown in Figure B.2.

Elevation, Proximity to Pacific Coastline and Housing Density

A number of factors may determine the level of air pollution in a neighborhood. For example, all other things equals, air pollution will generally be less in coastal communities because of the prevailing west winds. In addition, local climate conditions are likely to have a significant impact on the concentration of air pollutants. Also, densely populated urban areas generally tend to have more air pollution because of higher road congestion. To account for these factors we add three neighborhood variables to the household preference specification. These are the mean elevation of the neighborhood, the neighborhood's proximity to the Pacific coastline and the housing density of the neighborhood.

The National Elevation Dataset (NED) is a product of the US Geological Survey. It was developed by merging the highest resolution and best quality elevation data across the United States into a seamless raster format. The data are provided at a resolution of 1 arc second with the unit of elevation in meters. We use the NED to calculate the average elevation of each PUMA. The neighborhood's proximity to the Pacific coastline is measured by a binary variable which equals one if a portion of the neighborhood's boundary is on the Pacific coastline. The housing density of the PUMA is given by the number of housing units per square kilometer.

C. Asymptotic Properties of Parameter Estimates

Identification

We discuss the identification of the parameters of the household's indirect utility function. Specifically, we ask what features of the data allow for the identification of the estimated parameters. A separate, though not unrelated, identification argument can be given for the each of the stages of the estimation.

A necessary data requirement for identification of the first-stage parameters is that the number of observations be larger than the number of alternative-specific constants (H-1) plus the number of interaction parameters (k) to be estimated. In particular let N be the number of households in the sample. Then we must have that $N \geq H + k - 1$. Note that this condition has a direct implication for the characterization of residential locations and the household sample. First, it implies that the household sample used in the estimation must be at least of size $H + k - 1$. Second, characterizing the residential locations as individual housing units would imply that $N < H + k - 1$. As a result, the alternative constants may not be identified, hence the need to characterize residential locations using housing products rather than individual houses.

Given that the data satisfy the necessary requirement for identification, the heterogeneous taste (i.e., interaction) parameters will be identified, provided that there are sufficient differences in the attributes of households' location choices across each dimension of the household characteristics. For instance, suppose we hypothesize that college-educated households have a higher WTP for school quality relative to the remainder of the population. Then, for the interaction parameter between school quality and college education to be identified, we need to observe a sufficient difference (in this case positive) in the school quality levels of residential locations chosen by college-educated households compared to the remainder of the households.

The alternative-specific constants, which will characterize the mean utility from each residential location, are identified by the variation in the market shares[26] of residential locations. Simply put, if residential location A is on average preferred to residential location B (i.e.,

[26] The market share of a housing product is defined as the proportion of households choosing the housing product in the 1990 PUMS data.

$\delta_A > \delta_B$) then, all other things equal, we should observe more households choosing A over B in the data.

The mean taste parameters in the second-stage regression are identified from the variation in the market shares of residential locations across housing and neighborhood attributes. Notice that a necessary condition is that the alternative-specific constants are identified in the first-stage estimation. This should obviously be the case, since the second-stage regression cannot be defined without the alternative-specific constants. We can illustrate the second-stage identification argument as follows. Suppose, for example, that we hypothesize that households place, on average, a negative value on air pollution. Then in order to identify the negative mean taste parameter for air pollution we must observe that, holding all other attributes equal, residential locations in highly polluted areas have a lower market share compared to residential locations in the least-polluted areas.

Consistency and Asymptotic Normality

Similar to the identification argument, the asymptotic properties of the estimates can be discussed in terms of the first- and second-stage estimation. An in-depth discussion of the asymptotic properties of the two-stage estimator can be found in Bayer et al. (2005). The consistency and asymptotic normality of the first-stage estimates follow in the same spirit as in the traditional multinomial logit estimation. Given identification of the first stage, the estimated alternative-specific constants and heterogeneous taste parameters will be consistent and asymptotically normal as long as the number of households (N) in the sample grows large (Bayer et al., 2005).

The argument for consistency of the second stage is, however, less straightforward. The complication arises because the dependent variable in the second-stage regression is the estimated vector of alternative-specific constants from the first stage. Hence a large number of housing products is not sufficient to guarantee consistency and asymptotic normality. A formal proof is given in Berry, Linton and Pakes (2004). They show that the second-stage estimates will be consistent as long as (i) the number of housing alternatives, H, grows large and (ii) $H \log H/N$ goes to zero. That is, not only must H grow large but the number of households in the sample must also grow faster than H. In addition, asymptotic normality at a rate $\sqrt{H}$ requires that H^2/N be bounded. In other words, N must grow at a rate faster than H^2.

D. Additional Robustness Checks: Alternative Sampling Strategies

We check the robustness of the estimated preference parameters with respect to the size of the household's sampled choice set. In section 5.1.2 we explained that the household's relevant choice set includes the (i) chosen alternative and (ii) a random sample of 20 non-chosen alternatives. Model 5a in Table D.1 re-estimates the preference parameters using a choice set that includes (i) the chosen alternative and (ii) a random sample of 10 non-chosen alternatives. Model 5b uses a random sample of 50 non-chosen alternatives to form the household choice set. The estimated parameters from both specifications have the same signs with the coefficients in Model 1. The magnitudes of the estimated parameters are also very similar across the specifications.

We do a final robustness check of the estimated parameters with respect to the sampling of the households. In section 5.1.1 we explained that the household sample is formed by drawing a 10 percent random sample of the households choosing each housing type. We re-estimate the household parameters using a different sample size for the random draws. The results are reported in Models 6a and 6b of Table D.1. Model 6a reports the estimates from a household sample obtained by drawing 20 percent of the households choosing each housing type. Model 6b reports the estimates from a household sample obtained by drawing 40 percent of the households choosing each housing type. The estimated coefficients are also very similar to those in Model 1.

E. Comparing MWTP Estimates with the Existing Literature

Table E.1 summarizes the marginal willingness to pay (MWTP) estimates, in annual dollar terms, for selected housing and neighborhood characteristics. The mean MWTP, in annual dollar terms, for a housing attribute x_k is defined as:

$$MWTP_k = \frac{\overline{(\partial \hat{V}_{ih} / \partial x_k)} \cdot 12}{\hat{\alpha} / \overline{(y - p)}}. \tag{4}$$

Where, $\hat{V}_{ih}$ is the estimated household indirect utility function, and $\overline{\partial \hat{V}_{ih} / \partial x_k}$ represents the marginal utility of x_k evaluated at the mean of the household sample. The term in the denominator represents the marginal utility of income evaluated at the mean of the household sample. The mean MWTP for a specific group of the household population (i.e., college graduates, annual income below $19,000) is obtained by evaluating the marginal value and the marginal utility of income at the group mean.

All things equal, we find that households are willing to pay an additional $1,100 in annual housing rent for an extra bedroom in their house. Households are willing to pay an additional $10,000 annually or nearly twice the average annual rent to reside in a single-family housing unit. Households are also willing to pay an additional $9,800 annually, a one standard deviation increase in neighborhood school quality. The model also predicts that households will pay nearly twice the average rent to live in coastal communities. The estimated mean MWTP for locations that are within the household’s employment zone is very large. Households are, on average, willing to pay roughly six times the average annual rent for locations that are within their employment zone. As explained earlier, this may be due to the fact that the employment zone dummy may be capturing other unobserved neighborhood characteristics that are valued by households. MWTP estimates also vary across household characteristics. For instance, compared to the average household, college graduates will pay an extra $500 per year for a one-point increase in the neighborhood schools' average math score. Math scores range from 25 to 60 in the study area.

Our estimate of the MWTP for air quality ($62) compares well with other estimates in the literature. Sieg et al. (2004) report a MWTP of $61 for a 1 percent reduction in the 1990 average ozone concentration. Estimates of the MWTP for air quality range from $18 to $181 in the literature (Sieg et al., 2004). The estimates of the MWTP for bedrooms also vary in the literature. Bayer et al. (2005) find a mean MWTP of $1,312, in annual 1990 dollars, for an additional bedroom. Quigley (1985) estimates a nested logit model of household choice in the Pittsburgh metropolitan area, and finds that households are, on average, willing to pay $618 in annual 1990 dollars for an additional bedroom.[27] On the other hand, Chattopadhyay (2000) estimates a similar model for the Chicago area using four alternative nesting structures. He finds that the WTP for an additional bedroom ranges from $82 to $533, in 1990 annual dollars.

Our estimate of the mean MWTP for a one standard deviation increase in school quality is very large compared to the estimate obtained Bayer et al. (2005). Our mean MWTP estimate for a one standard deviation increase in the school quality level is $3,550 in annual terms. This compares with the Bayer et al.'s estimate of $21.5. It should be noted, however, that the two school quality measures are reported using different scales. The mean school quality in the Bayer

[27] The estimate reported in the paper is $13.18 per month in 1967 dollars. This estimate is converted into annual 1990 dollars.

et al. sample is 527, while our school quality measure has a mean of 34. As a result, it makes sense to also compare the mean MWTP for a 1 percent change in the annual 1990 mean school quality, as suggested by Sieg et al. (2004) in the case of air quality. Our estimate of the mean MWTP for a 1 percent change in the mean school quality level is $136, which is much closer to the Bayer et al. estimate of $18.

We would expect our MWTP estimate of to be relatively higher than the estimate from Bayer et al. (2005). This is because, in our model, school quality may be correlated with other unobserved neighborhood-quality characteristics contained in ξ_h. As result, the second-stage OLS regression may tend to overestimate the mean taste for school quality. Bayer et al. control for this problem using school district boundary fixed effects. It is not possible to apply this approach to our data because the neighborhoods, ie. PUMAs, are too large compared to school districts. The neighborhoods in Bayer et al. are Census blocks, which are much smaller geographic units compared to school districts. This facilitates the use of school district boundary fixed effects because most Census blocks fall within the boundary of a school district, whereas most PUMAs do not.

F. Simulation of the Counterfactual Equilibrium for the Year 2000

F1. Calibrating the Housing Demand

The economic agents in this model are households. We consider the housing choices of N_s (=17,894) households sampled from the overall population of N_p (=171,000) households obtained from the 1990 Census PUMS. The sampling framework used to generate the household sample is described in section 5.1. The housing market is characterized by 4,037 distinct housing types. The choice set of each sampled household is characterized by the sampling framework in section 5.1.

We could have each household facing the full set of 4,037 housing types. However, this would not be consistent with the estimation of household preference parameters. Recall that the maximum likelihood estimation, which uses choice set sampling, ensures that the market is in equilibrium in the 1990 benchmark (see section 5.2.1). This benchmark equilibrium, which is enforced via the first-order conditions of the maximum likelihood estimation (see equation 5.3), will no longer hold when households face the full set of alternatives.[28] As a result, significant errors arise in the computation of the predicted housing-type demands, and the counterfactual equilibrium housing-type prices may have undesirable[29] properties. Hence, we prefer to maintain the choice set sampling framework, used during the estimation, in the calibration of housing-type demands. We next discuss strategies for obtaining consistent estimates of housing-type demands under choice set sampling.

Obtaining an Appropriate Forecast of the Demand for Housing Types

The computation of the counterfactual equilibrium begins with forecasting the predicted demand for each housing type in the household population under the new air quality levels. Ben-Akiva and Lerman (1985) provide a detailed overview of various techniques for obtaining appropriate forecasts of aggregate demands for choice alternatives in discrete choice models. Our prediction of the aggregate demand for a residential location *h* uses the method of sample enumeration. This technique is especially appropriate in cases when (i) the household sample is drawn nonrandomly[30] from the population and (ii) the choice set of the household is formed by taking a random subsample of the full set of alternatives. In both of these cases sample enumeration

[28] In order for the benchmark equilibrium to hold we will need to re-estimate the preference parameters using the full choice set, which is not computationally feasible.
[29] Notably, some housing types may have negative prices in the counterfactual equilibrium.
[30] Our household sample is formed by drawing a 10 percent random sample of the households choosing each housing type.

allows the researcher to obtain a consistent[31] estimate of the share of the household population choosing a residential location *h*. For a stratified sample with $g = 1,\ldots, G$ strata, Ben-Akiva and Lerman (1985) define the sample enumeration estimate of the share of the household population choosing an alternative *h* as:

$$\hat{\sigma}_h = \sum_{g=1}^{G} \left(\frac{N_g}{N_p} \right) \frac{1}{N_{sg}} \sum_{i=1}^{N_{sg}} \hat{P}_{ih} (\hat{\delta}^{mle}, \hat{\alpha}, \hat{\gamma}, \hat{\beta}_1), \tag{5}$$

where, N_p is the household population, N_g is the population size of strata *g*, N_{sg} is the sample size of strata *g*, and $\hat{P}_{ih}$ is the estimated household choice probability. For the sampling design used in this study (see section 5.1.1), each housing type *h* represents a stratum. As a result $N_g = N_h$, $N_{sg} = N_{sh}$, the first summation term drops out and the expression for the estimated population share becomes:

$$\hat{\sigma}_h = \left(\frac{N_h}{N_p} \right) \frac{1}{N_{sh}} \sum_{i=1}^{N_{sh}} \hat{P}_{ih} (\hat{\delta}, \hat{\alpha}, \hat{\gamma}, \hat{\beta}_1), \tag{6}$$

Where, N_h is the number of households choosing location *h* in the population and N_{sh} the number of households choosing location *h* in the household sample. The population share is then converted into the predicted population demand for a housing location *h* by multiplying the estimated share by the household population (N_p). For a given housing type *h* the predicted population demand is given by:

[31] Consistency of the estimated population share holds as long as the estimated preference parameters are consistent, which is the case in our estimation.

$$\hat{d}_h(p) = \frac{N_h}{N_{sh}} \sum_{i=1}^{N_{sh}} \hat{P}_{ih}(\hat{\delta}, \hat{\alpha}, \hat{\gamma}, \hat{\beta}_1) \,. \tag{7}$$

The main limitation of the sample enumeration estimate, of the predicted population demand, is that it is subject to sampling error. The sampling error is due to the sampling of households and the sampling of the household choice sets. However, in our application, the sampling error is relatively small given the large size of our sample. The sampling error in the predicted population share for housing type *h* can be computed using the weighted root mean square formula provided by Ben-Akiva and Lerman (1985), which is due to Koppelman (1975). For our sampling framework, the sampling error in estimating the population shares for the 1990 benchmark is given by:

$$rms = \left\{ \sum_{h=1}^{H} \frac{N_h}{N_p} \hat{\sigma}_h \left[\frac{\sigma_h - \hat{\sigma}_h}{\sigma_h} \right]^2 \right\}^{1/2} , \tag{8}$$

where σ_h represents the actual share of the household population choosing housing type *h*, which in our sampling framework turns out to equal N_h / N_p. The weighted root mean square in our application is approximately 10^{-10} which is fairly small. An alternative way of assessing the sampling error is to compute the square root of the sum of squares of the excess demands across housing types in the benchmark. This is because, by virtue of the maximum likelihood estimation, the benchmark excess demands[32] must equal zero if there is no sampling error in the

[32] The excess demands are given by $\hat{d}_h - s_h$, where s_h is the supply of housing units of type *h*.

predicted population demand. The sampling error in the predicted population demand can then be obtained as $\left\{ \sum_{h=1}^{H} \left(\hat{d}_h - s_h \right)^2 \right\}^{1/2}$. In our application the sampling error in the predicted population demand is roughly 10^{-6} which is also small.

Computing the Predicted Population Demand under the New Air Quality Levels

Using equation (7) we can now characterize the predicted population demand for each housing type under the new air quality levels. It is given by:

$$\hat{d}_h^1(p) = \frac{N_h}{N_{sh}} \sum_{i=1}^{N_{sh}} \frac{\exp[\hat{\delta}_h + \hat{\alpha}\log(y_i - p_h) + \hat{\gamma} d_{ih} + \sum_k x_{hk}^1 z_{ir} \hat{\beta}_{1kr}]}{\sum_{m \in C_i} \exp[\hat{\delta}_m^1 + \hat{\alpha}\log(y_i - p_m) + \hat{\gamma} d_{ih} + \sum_k x_{mk}^1 z_{ir} \hat{\beta}_{1kr}]}, \qquad (9)$$

where, $\hat{P}_{ih}$ has been defined explicitly. C_i represents the choice set of household i. The superscript 1 is used to indicate market conditions after the air quality changes have occurred. x_{kh}^1 is the vector of attributes for housing type h which includes the new air quality level. $\hat{\delta}_h^1$ represents the predicted base utility for housing type h under the new air quality levels. It is given by:

$$\hat{\delta}_h^1 = \sum_k x_{hk}^1 \hat{\beta}_{0k} + \hat{\xi}_h,$$

where $\hat{\xi}_h$ is the vector of residuals obtained in the second-stage OLS estimation (Equation 5.7). $\hat{\xi}_h$ characterizes the estimate of the mean valuation from the unobserved location attributes. The vector of residuals must be added because the alternative constants which characterize the benchmark 1990 equilibrium are given by:

$$\delta_h^0 = \sum_k x_{hk}^0 \beta_{0k} + \xi_h \,.$$

The reader can note that this is the same equation characterizing the mean utility in equation (5.4). Hence ξ_h is an key component of the functional form of δ_h. In our application, the second-stage taste parameter for air quality is not statistically different from zero. Hence the predicted base utility levels ($\hat{\delta}_h^1$) under the new air quality levels are the same as the benchmark base utility levels ($\hat{\delta}_h^0$) obtained from the first-stage of the estimation.

F2. Defining the Counterfactual Locational Equilibrium

The 171,000 housing units occupied by the population of households in the 1990 Census PUMS are classified into 4,037 residential locations. The housing supply s_h is given by the number of housing units at each residential location h. We assume that the housing supply is exogenous with respect to the changes in air quality. Given the housing supply (s_h) and the predicted housing demand ($\hat{d}_h^1$), the counterfactual equilibrium price vector is defined by:

$$ed_h(p^*) = \hat{d}_h^1(p^*) - s_h = 0 \qquad h = 1, \ldots, H. \tag{10}$$

The counterfactual locational equilibrium defined by equation (10) is unique and locally stable. This follows from the fact that the parameter estimate $\hat{\alpha}$ is positive and hence the excess demand $ed_h(p)$ satisfies the strict gross substitution property (see Anas, 1982).

F3. Implementation

A numerical solution to the system of H equations in H variables, which defines the counterfactual locational equilibrium, is obtained via an efficiently convergent algorithm suggested by Anas (1982). The equilibrium price vector is found iteratively via a price-adjustment process that starts with the benchmark 1990 price vector p^0 and adjusts the location prices until the adjusted price vector is arbitrarily close to the equilibrium price vector p^*.

Let $t = 1, \ldots, T$ define a sequence of T iterations such that $p^T \approx p^*$. The price vector at iteration $t + 1$ is given by the Newton step:

$$p^{t+1} = p^t - [\partial ed(p^t)/\partial p]^{-1}[ed(p^t)]. \tag{11}$$

$ed(p)$ represents the system of excess demands for all H residential locations, and $[\partial ed(p^t)/\partial p]$ is the Jacobian matrix of $ed(p)$ evaluated at p^t. Computation of the Newton step defined in (11) requires evaluating and inverting the Jacobian which has dimension $H = 4037$. The computational cost of this algorithm is considerably large. The evaluation of the Jacobian alone takes approximately 30 minutes on a Pentium 4 2Ghz PC station.

Anas (1982) suggests a less costly iteration step which is obtained by ignoring the off-diagonal element of the Jacobian matrix. In this case the iteration step $t + 1$ is defined independently for each residential location h as:

$$p_h^{t+1} = p_h^t - ed_h(p^t)/[\partial ed_h(p^t)/\partial p], \qquad h = 1, \ldots, H. \tag{11a}$$

The computational cost of the iteration step in (11a) is significantly less than that of (11) since it only requires computing the diagonal vector of the Jacobian matrix and its element inverse. This alternate Newton step will converge to the equilibrium price vector p^* as long as the off-diagonal elements of the Jacobian are significantly small in absolute value compared to the diagonal elements. Convergence is achieved when the price vector p^T at iteration T is "sufficiently" close to p^*. In our counterfactual simulation p^T is considered "sufficiently" close to p^* if

$$ed_h(p^T)/s_h \leq 10^{-5} \qquad h = 1,\ldots, H. \qquad (12)$$

In other words, the absolute absolute value of the excess demand for each location is less than 0.001% of the housing supply.

A computational issue arises from the fact that knowledge of H-1 housing-type excess demands is sufficient to characterize the system of H excess demands. This is because we assume that the housing market is a closed economy, which implies that no household relocates outside the study area. A direct implication of the closedness assumption is that the housing-type demands always sum to the total population (N) of households. This means that the system of H housing-type excess demands has only H-1 degrees of freedom. As a result, we fix one of the prices when solving for the numerical solution. This normalization guarantees that any starting value will lead to the same market clearing prices. The normalization also guarantees that the counterfactual equilibrium prices are within the same H-dimensional simplex as the benchmark price vector and hence lies in the positive quadrant $\Re^{H+}$.

References

Anas, A. 1982. *Residential Location Markets and Urban Transportation: Economic Theory, Econometrics and Policy Analysis with Discrete Choice Models*. Academic Press, New York.

Bayer, P., R. McMillan and K. Rueben. 2005. "An Equilibrium Model of Sorting in an Urban Market." NBER Working Paper: 10865.

Ben-Akiva, M., and S. R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge.

Berry, S., O. B. Linton and A. Pakes. 2004. "Limit Theorems for Estimating the Parameters of Differentiated Product Demand Systems." *Review of Economic Studies* 71(3):613-54.

Chattopadhyay, S. 2000. "Effectiveness of McFadden's Nested Logit Model in Valuing Amenity Improvements." *Regional Science and Urban Economics* 30(1): 23-43.

Koppelman, F. 1975. *Travel Prediction with Models of Individualistic Choice Behavior*. Ph.D. dissertation, Department of Civil Engineering, MIT, Cambridge, Mass.

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

Quigley, J. M. 1985. "Consumer Choice of Dwelling, Neighborhood and Public Services." *Regional Science and Urban Economics* 15(1):41-63.

Sieg, H., V. K. Smith, H. S. Banzhaf and R. Walsh. 2004. "Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods." *International Economic Review* 45(4): 1047-77.

Table A1. Regression Used for Correcting House Values

	Los Angeles – Long Beach MSA	Orange County MSA	Riverside – San Bernardino MSA
Log Transaction price (10 times property tax)	.335**	.349**	.440**
Moved in 1985 to 1988 (compared to 1989-90)	-0.013**	0.017**	-0.060**
Moved in 1980 to 1984	0.037**	0.075**	-0.071**
Moved in 1970 to 1979	0.192**	0.309**	0.040**
Moved in 1960 to 1969	0.253**	0.395**	0.097**
Moved in 1959 or earlier	0.201**	0.307**	0.088**
R^2	0.325	0.263	0.431
Observations	138,181	39,550	33,891

Note: * Significant at the 5 percent level. ** Significant at the 1 percent level. Dependent variable is log of house value. Regression includes a full set of PUMA dummies.

Table A2. Regression Used for Correcting Monthly Rents

	Los Angeles – Long Beach MSA	Orange County MSA	Riverside – San Bernardino MSA
Moved in 1985 to 1988 (compared to 1989-90)	-0.082**	-0.062**	-0.081**
Moved in 1980 to 1984	-0.207**	-0.193**	-0.234**
Moved in 1970 to 1979	-0.329**	-0.298**	-0.328**
Moved in 1960 to 1969	-0.410**	-0.439**	-0.295**
Moved in 1959 or earlier	-0.421**	-0.310**	-0.459**
Rooms	0.027**	0.014**	0.043**
Bedrooms	0.154**	0.144**	0.121**
Single-family house attached (compared to single-family detached)	-0.056**	-0.029**	-0.080**
2 apartments complex	-0.098**	-0.128**	-0.182**
3-4 apartments complex	-0.128**	-0.132**	-0.168**
5-9 apartments complex	-0.144**	-0.168**	-0.174**
10-19 apartments complex	-0.142**	-0.166**	-0.133**
20-49 apartments complex	-0.113**	-0.138**	-0.143**
50 or more apartments complexes	-0.145**	-0.170**	-0.148**
Built in 1985 to 1988 (compared to 1989-90)	-0.001**	0.011**	-0.053**
Built in 1980 to 1984	-0.089**	-0.073**	-0.139**
Built in 1970 to 1979	-0.078**	-0.045**	-0.158**
Built in 1960 to 1969	-0.068**	-0.072**	-0.211**
Built in 1950 to 1959	-0.105**	-0.103**	-0.257**
Built in 1940 to 1949	-0.122**	-0.149**	-0.291**
Built in 1939 or earlier	-0.146**	-0.161**	-0.322**
R^2	0.368	0.410	0.395
Observations	138,181	39,550	33,891

Note: * Significant at the 5 percent level. ** Significant at the 1 percent level. Dependent variable is monthly rent. Regression includes a full set of PUMA dummies.

Table A3. Regression Used for Converting House Values to Rental Rates

	Los Angeles – Long Beach MSA	Orange County MSA	Rverside – San Bernardino MSA
Owner-occupied	5.654**	5.830**	5.474**
Rooms	0.049**	0.042**	0.090**
Bedrooms	0.052**	0.080**	0.036**
Single-family house attached (compared to single-family detached)	-0.126**	-0.080**	-0.015**
2 apartments complex	-0.217**	-0.143**	-0.160**
3-4 apartments complex	-0.210**	-0.179**	-0.130**
5-9 apartments complex	-0.232**	-0.207**	-0.138**
10-19 apartments complex	-0.229**	-0.204**	-0.108**
20-49 apartments complex	-0.201**	-0.166**	-0.111**
50 or more apartments complexes	-0.234**	-0.188**	-0.119**
Built in 1985 to 1988 (compared to 1989-90)	-0.014**	-0.063**	-0.049**
Built in 1980 to 1984	-0.083**	-0.158**	-0.142**
Built in 1970 to 1979	-0.105**	-0.191**	-0.200**
Built in 1960 to 1969	-0.182**	-0.239**	-0.303**
Built in 1950 to 1959	-0.247**	-0.293**	-0.382**
Built in 1940 to 1949	-0.255**	-0.316**	-0.402**
Built in 1939 or earlier	-0.257**	-0.319**	-0.409**
R^2	0.992	0.987	0.986
Observations	138,181	39,550	33,891

Note: * Significant at the 5 percent level. ** Significant at the 1 percent level. Dependent variable is log of corrected house value if owned, otherwise, log of corrected monthly rent. Regression includes a full set of PUMA dummies.

Table D.1: Alternative Sampling Strategies

	Model 1‡	Model 5a	Model 5b	Model 6a	Model 6b
First Stage					
Log(y-p)	1.475**	1.394**	1.536**	1.6**	1.56**
Ozone * Log(y-p)	-0.019**	-0.023**	-0.019**	-0.022**	-0.021**
Bedrooms * Household size	0.066**	0.07**	0.066**	0.062**	0.053**
Single family * Children under 18	0.227**	0.271**	0.225**	0.21**	0.253**
Math * College educated head	0.309**	0.32**	0.3**	0.295**	0.297**
Log crime * Log(y-p)	0.004**	0.001	0.006**	0.006**	0.009**
Within household's employment zone	1.989**	1.986**	2.006**	1.971**	1.961**
Log-Likelihood	-37,072	-27,104	-51,690	-67,241	-130,056
Likelihood Ratio p-value (H_0: δ =0)	0.000	0.000	0.000	0.000	0.000
McFadden pseudo-R^2	0.319	0.368	0.265	0.353	0.365
Observations	17,894	17,894	17,894	34,132	67,304
Second Stage OLS †					
Bedrooms	0.04*	0.03	0.045*	0.048*	0.05**
Built after 1980	-0.594**	-0.602**	-0.588**	-0.596**	-0.59**
Built in 60s or 70s	-0.172*	-0.18**	-0.168*	-0.18**	-0.175**
Single-family dwelling	0.352**	0.359**	0.356**	0.351**	0.355**
Owned	0.054	0.06	0.041	0.045	0.047
Math test score	0.139**	0.13**	0.143**	0.147**	0.138**
Log FBI crime index	0.0005	0.001	0.0001	0.000	0.000
Log elevation	0.016	0.009	0.025	0.028	0.028
PUMA is on pacific coastline	0.342**	0.341**	0.334**	0.341**	0.349**
Log density	0.079	0.07	0.089*	0.09*	0.094*
Prop. of population Hispanic	-0.38*	-0.401*	-0.387*	-0.376*	-0.41**
Ozone	0.161	0.19	0.148	0.151	0.141
R^2	0.054	0.052	0.055	0.056	0.057
Observations	4,037	4,037	4,037	4,037	4,037

Notes:
** Significant at 1% level. * Significant at 5% level. † Standard errors are computed using White's robust covariance matrix.
‡ Model 1 : Benchmark specification used in the simulation and welfare analysis.
Model 2a: Characterizes the household's relevant choice set using 10, instead of 20, randomly sampled non-chosen alternatives.
Model 2b: Characterizes the household's relevant choice set using 50, instead of 20, randomly sampled non-chosen alternatives.
Model 3a: household sample is form by drawing 20, instead of 10, percent of the households choosing each alternative in the 1990 PUMS.
Model 3b: household sample is form by drawing 40, instead of 10, percent of the households choosing each alternative in the 1990 PUMS.

Table E.1: MWTP for Selected Housing and Neighborhood Attributes (1990 Annual Dollars)

	Mean MWTP	College Grads	Income < \$19,600	Income > \$60,400
Bedrooms (+1 bedroom)	1,143	-	150	2,426
Single-family dwelling (vs. Multi-family)	10,104	-	1,326	21,450
Math test score* (+1 standard deviation)	3,550	11,474	466	7,538
PUMA is within household's employment zone (vs. outside)	57,119	-	7,494	121,262
PUMA is on Pacific coastline (vs. inland)	9,821	-	1,289	20,850
Share Hispanics (+0.01)	-109	-	-14	-232
Ozone† (-1%)	62	-	8	131

Note: All values are in annual rental rates. For example, the average household is willing to pay \$1,143 annually for an additional bedroom whereas households with income below \$19,000 are only willing to \$150. The annual mean rental housing price in the study area is \$9,000.

†MWTP for a 1 percent change in 1990 average.

* Math test score: mean = 34, standard deviation = 8.9, range: 25 to 60.

Figure B.1: 1990 Neighborhood School Quality Levels

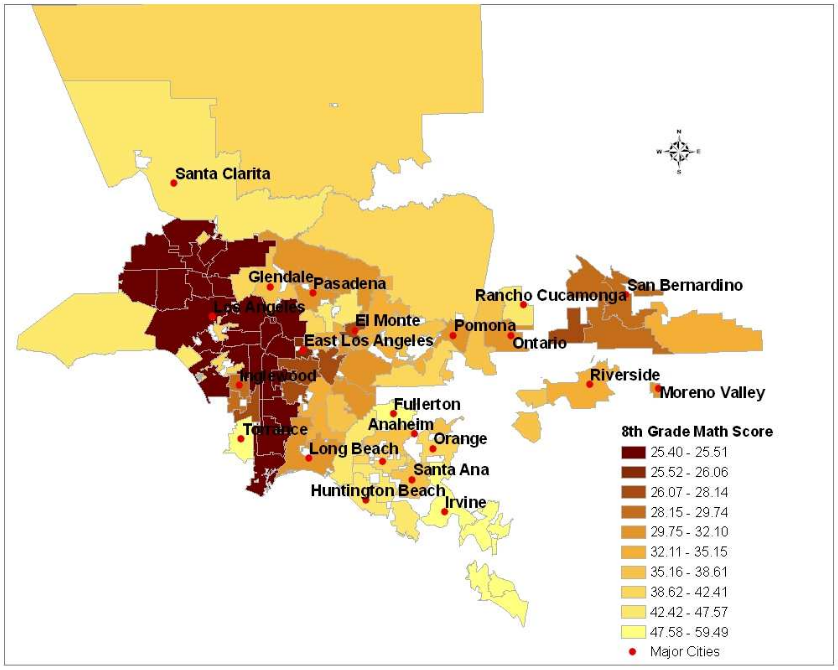

Figure B.2: 1990 Neighborhood Crime Levels