

Decomposing the Structural Identification of Non-Market Values

Forthcoming: *Journal of Environmental Economics and Management*

Nicolai V. Kuminoff *

Department of Economics

Arizona State University

kuminoff@asu.edu

August 2009

* I thank V. Kerry Smith, Paul Fackler, Roger von Haefen, Arik Levinson, Ray Palmquist, Dan Phaneuf, Jaren Pope, Chris Timmins, and two anonymous referees for helpful comments and suggestions. I also thank seminar participants at the National Center for Environmental Economics, North Carolina State University, University of California, Berkeley, University of Minho (Portugal), University of Wyoming, University of Tennessee, and Virginia Tech. Support from the National Science Foundation and the Virginia Agricultural Experiment Station is gratefully acknowledged.

Decomposing the Structural Identification of Non-Market Values

ABSTRACT:

This paper addresses the issue of transparency in structural modeling by investigating how the assumptions of a structural model contribute to its identification of nonmarket values.

Assumptions about the form of the utility function, the depiction of preference heterogeneity, and the definition for the choice set are shown to jointly identify a partition of preference space which explains observed behavior. The borders of the partition reflect substitution patterns. The regions in the partition set-identify preferences. Assumptions about the distribution of preferences within each region point-identify partial equilibrium welfare measures. Together, these observations provide a framework for decomposing the structural identification of nonmarket values. The new framework allows analysts to explain how their modeling assumptions influence estimates for benefits and costs. This addresses federal guidelines for quantifying uncertainty in the modeling assumptions which enter benefit-cost analysis. The framework is demonstrated by decomposing estimates of the MWTP for air quality.

KEY WORDS: Equilibrium sorting, locational equilibrium, partition, hedonic, random-utility, willingness-to-pay, preferences.

JEL CODES: C35, Q51, R21, H41

1. Introduction

People differ in their preferences for amenities, they interact with their neighbors, and they adjust their behavior in response to large changes in health risk and environmental quality. Recent advances in discrete choice modeling have allowed us to incorporate these basic notions about consumer behavior into models of the demand for nonmarket goods and services. While the new models can produce benefit estimates which reflect preference heterogeneity, social interactions, and general equilibrium adjustment, their increasing econometric complexity can make it difficult to understand how welfare measures are actually identified. This tradeoff is exemplified by equilibrium sorting models of household location choice.

Equilibrium sorting models build on the theory and econometrics of hedonic and discrete-choice models of differentiated product markets. They combine the information provided by an *equilibrium* hedonic price function [9, 21] with a formal description for the choice process that explains how heterogeneous consumers *sort* themselves across segments of a market [7, 17]. When applied to the market for housing, these models use information on households and their location choices to estimate structural parameters which characterize heterogeneity in preferences for local public goods and environmental amenities [6, 10]. The results can be used to predict the distributional welfare implications of future policy changes, while recognizing that a sufficiently large shock will induce some households to move [24, 25]. Equally important, these models are capable of reflecting feedback effects between people and their surrounding environment. For example, while people are affected by environmental amenities, their actions can degrade the ecosystems providing those amenities. The ability to model these interactions and analyze their welfare implications has expanded the scope for public policy analysis.

Recent applications have analyzed the welfare implications of large scale changes to air

quality, school quality, open space, climate, and other amenities [5, 11, 24, 25, 26, 30]. These studies demonstrate that preference heterogeneity and feedback effects can have first-order policy implications. For example, Smith et al. [25] find that the benefits from air quality improvements attributed to the 1990 Clean Air Act Amendments vary dramatically within the Los Angeles region—from \$33 to \$2400 per household. This range reflects differences in initial air quality conditions, income, and general equilibrium price adjustments that occur as some households react to improved air quality by moving. Similarly, Walsh’s [30] analysis of the demand for open space in North Carolina illustrates how the migration patterns which follow major land use changes can stimulate additional urban development. This feedback effect dominates welfare measures and helps to explain long-run development patterns.

The potential for using equilibrium sorting models to conduct high resolution policy analysis is exciting, but are their predictions reliable? In recent years, economists have become increasingly skeptical of structural modeling. While structural models require us to formalize our assumptions about the choice set, utility functions, and random parameters, the extent to which these assumptions ultimately drive empirical results is often unclear. Some researchers have turned to experimental and quasi-experimental methods as a way to achieve more transparent identification. Others have called on structural modelers to clarify how their assumptions contribute to the identification of key results [12, 13]. Achieving this clarity is especially important in benefit-cost analysis, where structural modeling has the potential to influence environmental policy. This is underscored by the Office of Management and Budget’s guidelines for benefit-cost analysis of federal programs, which requires economists to quantify how their modeling assumptions influence estimates for benefits and costs [28].¹

This paper addresses the issue of transparency in structural modeling by presenting a new

framework for decomposing the identification of non-market values. It uses an equilibrium sorting model of household location choice to illustrate how our assumptions about choice sets, utility functions, and random parameters influence what we learn about the willingness-to-pay for policy changes. While equilibrium sorting provides a timely example, the framework for decomposition is not limited to these models or to housing market applications. Its key features apply generally to structural discrete-choice models of the demand for a differentiated product.

The new framework builds on an extensive literature evaluating the role of structural assumptions in discrete-choice models (e.g. [18, 20, 27]). It begins by extending Varian's [29] diagrammatic analysis of recoverability to illustrate how heterogeneous preference parameters are jointly identified by structural assumptions and observed behavior. This illustration leads to four insights that are broadly relevant within the literature. First, increasing the degree of preference heterogeneity increases the sensitivity of welfare measures to distributional assumptions. Second, welfare measures can be decoupled from assumptions about the distribution of preferences. Third, ignoring moving costs and other sources of friction in the market creates a false sense of precision in estimates for welfare measures. Finally, the identifying power of assumptions about the choice set and consumer mobility can be distinguished from the identifying power of assumptions on preferences. After describing these results, the new framework is used to decompose estimates of the willingness-to-pay for improved air quality from a model of equilibrium sorting in Northern California. Given a parametric form for the utility function and a definition for the choice set, estimates for *average* benefits can vary by up to 167% depending on the econometrician's assumption about the distribution of preferences. However, for a subset of *individual* households, benefit estimates can be identified to within a few percentage points, regardless of distributional assumptions.

The remainder of the paper is organized as follows. Section two provides background on equilibrium sorting models of household location choice by describing how the sorting process relates to the hedonic price function which characterizes equilibrium in the housing market. Translating data on equilibrium outcomes into structural preference parameters requires assumptions on: (1) the shape of utility functions, (2) the form of preference heterogeneity, and (3) the distributions used to characterize this heterogeneity. Section three illustrates how these assumptions interact with the data and the definition for the choice set to identify welfare measures for policy changes. Section four uses this framework to decompose the identification of the Northern California model. After interpreting the results, section five concludes.

2. The Household's Location Choice Problem

Hedonic and equilibrium sorting models of household location choice typically start from the following problem: the availability of housing and public goods varies across an urban landscape and each household chooses to occupy the location in that landscape that provides its preferred bundle of goods, given its preferences, income, and the relative prices involved. Households are assumed to be freely mobile and each household pays for its location choice through the price of housing. In order to link a household's observed choice to its demand for an individual public good, the problem can be formalized using the characteristics approach to consumer theory [15]. This keeps the dimensionality of the problem manageable and allows locations to be treated as a differentiated product.

2.1. Characterizing Behavior in a Locational Equilibrium

Let the urban landscape consist of $j = 1, \dots, J$ locations, each of which is defined by a

bundle of housing characteristics and public goods, h_j, g_j . The first element of the bundle, h_j , represents a vector of structural characteristics that fully describe the private good characteristics of housing at location j . This includes the number of bedrooms, the number of bathrooms, square feet, and lot size, to name a few. The second element in the bundle, g_j , is a vector of local public goods which are consumed exclusively by the occupants of location j . Examples include crime, school quality, air quality, and open space. Notice that “public goods” are defined here to include environmental amenities in addition to services produced from local taxes.

A household’s utility depends on the characteristics of housing and local public goods at its location and on its consumption of a composite numeraire private good, b . Households are heterogeneous. They differ in their income, y , and in their preferences, α . Let the population of households be indexed from $i = 1, \dots, I$. Then the utility obtained by household i at location j can be represented as: $U(b, h_j, g_j; \alpha_i)$. Each household is assumed to choose a location and a quantity of b that maximize its utility subject to a budget constraint:

$$\max_{j,b} U(b, h_j, g_j; \alpha_i) \text{ subject to } y_i = b + \tilde{p}_j. \quad (1)$$

In the budget constraint, the price of the numeraire is normalized to one, and \tilde{p}_j represents annualized expenditures on housing at location j . In other words, \tilde{p}_j is the cost of occupying a single home at location j for one year.

A locational equilibrium is achieved when every household occupies its utility-maximizing location and nobody wants to move, given housing prices, housing characteristics, and the provision of local public goods. Under mild regularity conditions on $U(\cdot)$, the equilibrium relationship between \tilde{p}_j and h_j, g_j can be described by a hedonic price function:

$\tilde{p}_j = \tilde{p}(h_j, g_j)$. Specifically, if utility is continuously differentiable, monotonic in the numeraire, and Lipschitz continuous in h_j, g_j , it can be shown that equilibrium prices will be a function of all the choice-specific characteristics (see theorem 1 in [2]). This existence result does not rely on supply-side assumptions. Nor does it require continuity in the choice set. A hedonic price function will exist even if households choose among a discrete set of locations.² However, what we stand to learn from the hedonic price function in a locational equilibrium depends critically on whether the choice set is discrete or continuous.

2.2. Discrete and Continuous Choice Sets

Rosen [21] analyzed the properties of the hedonic price function that arises when heterogeneous consumers choose among a continuum of differentiated products. One of the key results from his analysis is that the partial derivatives of the price function provide a theoretically consistent measure of consumers' willingness-to-pay for marginal changes in product attributes. This result can be seen from the first order conditions to the utility maximizing problem.

$$\frac{\partial \tilde{p}(h_j, g_j)}{\partial g_j} = \frac{\partial U_{ij}(\cdot) / \partial g_j}{\partial U_{ij}(\cdot) / \partial b}, \quad \frac{\partial \tilde{p}(h_j, g_j)}{\partial h_j} = \frac{\partial U_{ij}(\cdot) / \partial h_j}{\partial U_{ij}(\cdot) / \partial b}. \quad (2)$$

Each household will maximize its utility by choosing a home that provides levels of each housing characteristic and public good at which the household's willingness-to-pay for an additional unit exactly equals its marginal implicit price. This result is frequently invoked to translate econometric estimates for $\tilde{p}(h_j, g_j)$ into households' willingness-to-pay for marginal changes in environmental amenities. The resulting estimates for the marginal willingness-to-pay can serve as an input into a "second-stage" estimation of the demand for an amenity (e.g. [19]). Furthermore, if the econometrician is willing to select a parametric form for the utility function,

equation (2) can be inverted to recover α directly, which can then be used to calculate households' ex ante willingness-to-pay for non-marginal changes in an amenity (e.g. [3]).

An important limitation of approaches to non-market valuation based on (2) is their reliance on continuity. If households are not free to choose continuous quantities of each structural characteristic and public good, equation (2) does not characterize equilibrium behavior. From an empirical perspective, this is a serious issue because many amenities vary across space in ways that are discrete. School quality varies discretely across school districts. Access to public pools, tennis courts and community centers may be limited to homeowners in a residential subdivision. Even for amenities which vary continuously in the natural environment, such as air quality and water quality, most of the variation tends to occur across natural boundaries, such as air basins and watersheds. Equilibrium sorting models provide a theoretically consistent way to recover α when one or more amenities cannot be chosen continuously.

Equilibrium sorting models depict households choosing among a finite number of residential communities, each of which provides a unique bundle of public goods. Public goods are assumed to be constant within each community, while the structural characteristics of housing vary continuously. Treating a residential community as a unique location implies an aggregation over individual homes. To address this aggregation, let p_j denote the annualized *per-unit* price of housing in community j and let q denote the quantity of housing consumed. Housing is now treated as a homogeneous commodity that can be consumed in continuous quantities. Under this assumption, the price of housing reflects the cost of consuming the public goods provided by each community.

Of course, in practice housing is not homogenous. Its structural characteristics vary

within and between communities, and these differences will be reflected in observable sale prices. This can be addressed if we are prepared to assume that the structural characteristics of housing enter the utility function through a sub-function that is homogeneous of degree one and separable from the effect of public goods and the numeraire. Under this restriction, Sieg et al. [23] demonstrate that the equilibrium hedonic price function will be separable in the structural characteristics of houses and the effect of public goods, as shown in (3).

$$\tilde{p}(g_j, h_{j,n}) = q(h_{j,n}) \cdot p_j(g_j). \quad (3)$$

The left side of the expression represents expenditures on house n in community j . The first term on the right side is a “quantity” index of housing that depends on the vector of structural characteristics $(h_{j,n})$. By condensing all the information about the structural characteristics into a single number, the index provides an empirical analog to the concept of a homogeneous unit of housing. The second term represents the price of a homogeneous unit of housing in community j , which depends only on the public goods it provides.³ Using this result together with the assumption that households are free to purchase any quantity of housing at the market price in each community allows preferences to be restated using the indirect utility function in (4).

$$V(g, p, \alpha, y) = U[g, q(g, p, \alpha, y), y - p \cdot q(g, p, \alpha, y), \alpha]. \quad (4)$$

Each household will choose the community that maximizes its utility, given income, prices, and the provision of public goods.

The indirect utility function in (4) describes how consumers choose among a discrete set of alternatives based on their income and preferences for the characteristics which differentiate those alternatives. This basic framework also forms the basis for random utility models (e.g. [27]), structural hedonic models (e.g. [3]), and corner solution models (e.g. [20]). All of these methods make parametric assumptions about $V(\cdot)$ in order to translate the observable

characteristics of consumers and their choices into structural preference parameters. The next section illustrates how these assumptions influence the resulting estimates for welfare measures and their policy implications.

3. Decomposing the Identification of Non-Market Values

When a consumer purchases a bundle of goods, the composition of that bundle reveals some basic features of his preferences. He could have chosen any other bundle within his budget. The fact that he did not reveals that he considers them to be inferior. Samuelson [22] demonstrated that comparing a bundle that was purchased with bundles that could have been purchased, but were not, can serve to identify bounds on a consumer's indifference curves.

To apply Samuelson's basic revealed preference logic to the location choice problem, suppose households choose among the four communities in Table 1. Each community differs in its price of housing and in its provision of two public goods, g_1 and g_2 . Consider a household in community C . If the household's preferences are monotonic and convex, the choice not to live in community B identifies bounds on the shape of that household's indifference curves. Figure 1 illustrates this graphically, using the technique from Varian [29]. The two points show the locations of communities B and C in public goods space. The shaded areas bound the set of indifference curves passing through point B that are consistent with a household that maximizes its utility by choosing to live in community C . For example, since preferences are convex and monotonic, any point on the line segment connecting B and C must provide at least as much utility as B , and any point to the northeast must be preferred to B . By the same reasoning, B must provide at least as much utility as any point on the line segment which connects B to the g_1 axis, and any point to the southwest must provide less utility than B . Repeating this exercise to

compare communities C and A would provide another set of bounds. While these bounds provide some intuition for the shape of a household's indifference curves in public goods space, they do not reveal much about willingness-to-pay for changes in those public goods.

In order to identify a household's indifference curve within the nonparametric bounds provided by general axioms on preferences, the econometrician must impose more structure on the problem. Two types of structural restrictions are required. First, a parametric indirect utility function must be selected. Second, a distribution must be specified for each preference parameter in that function used to characterize unobserved heterogeneity. Each restriction makes a different type of contribution to the identification.

Distributional assumptions are necessary due to the discreteness in the choice set. When household i chooses j from a finite set of communities, utility maximization is characterized by the set of inequalities in (5) rather than the first order conditions in (2).

$$V(g_j, p_j, \alpha_i, y_i) \geq V(g_k, p_k, \alpha_i, y_i), \quad \forall k = 1, \dots, J. \quad (5)$$

The equation simply says that if household i chooses location j it is because that location provides it with at least as much utility as any other location. Given a parametric form for the indirect utility function, the inequalities serve to identify a set of values for the heterogeneous parameters that are consistent with each observed choice. It must be the case that $(\alpha_i, y_i) \in \Omega_j$, where $\Omega_j = \{ (\alpha_i, y_i) : (\alpha_i, y_i) \text{ satisfies (5)} \}$. In words, the choice of community j reveals only that household i 's income and preferences lie somewhere in the Ω_j set. Imposing a distribution on α allows the econometrician to identify the density of preferences within Ω_j , given income.

To illustrate the role of each type of restriction in identifying preferences, suppose that utility can be represented by the following CES function:

$$V_{ij}(g_j, p_j, \alpha_i, y_i) = \left\{ \left[\exp\left(\frac{y_i^{1-\nu} - 1}{1-\nu}\right) \exp\left(-\frac{\beta(p_j^{\eta+1}) - 1}{1+\eta}\right) \right]^\rho + \alpha_{i0}(\bar{g}_{ij})^\rho \right\}^{\frac{1}{\rho}}, \quad (6)$$

$$\text{where } \bar{g}_{ij} = \alpha_{i1}g_{1j} + \alpha_{i2}g_{2j} \dots + \alpha_{iN}g_{Nj}.$$

This parameterization nests specifications used in previous studies to estimate households' preferences for school quality, crime rates, air quality, open space, and other amenities (e.g. [10, 24, 31]). The first term represents utility from the private good component of housing. β is an intercept in the demand for housing, and η and ν are price and income elasticities. This can be seen by using Roy's Identity to derive the housing demand function in (7).

$$q_i = \beta p_j^\eta y_i^\nu, \quad \beta > 0. \quad (7)$$

While the preference parameters that characterize this function are assumed to be constant across households, the demand still varies with income.

The second term in the CES function represents utility from public goods. \bar{g}_{ij} is a linear index of public goods provided by community j , as perceived by household i . Households differ in the relative weights they place on each public good in the index ($\alpha_{i1}, \dots, \alpha_{iN}$) and in the overall strength of their preferences for public goods relative to private goods (α_{i0}). Notice that α_{i0} is not identified separately from $\alpha_{i1}, \dots, \alpha_{iN}$. To distinguish their effects the weights can be normalized to sum to 1, allowing α_{i0} to be identified separately as a scaling parameter on the strength of preferences. This decomposition provides a convenient way to illustrate how alternative depictions of preference heterogeneity influence set identification of preferences.

3.1. Preference Heterogeneity, Set Identification, and Substitution Patterns

First consider the simplest form of preference heterogeneity where all the variation in

tastes can be condensed into a single heterogeneous preference parameter that ranks choices by their overall quality (see [10], [24], or [30] for examples). This is often described as “vertical differentiation”. The CES utility function in (6) simplifies to this case when households are constrained to have the same relative preferences for public goods: $\alpha_{i1}, \dots, \alpha_{iN} \equiv \alpha_1, \dots, \alpha_N, \forall i$.⁴

To see how vertical differentiation provides set identification of preferences, suppose households choose among the 4 communities in Table 1 based on the following relative preferences: $\alpha_{i1}, \alpha_{i2} \equiv 0.46, 0.54, \forall i$. With constant weights, all households agree on a common ranking of communities by the public goods they provide: $\bar{g}_{ij} \equiv \bar{g}_j, \forall i$.⁵ From most preferred to least preferred, the universal ranking is D, C, B, A . Now, given values for the housing demand parameters in (6), the system in (5) can be solved for bounds on the set of values for α_{i0} that rationalize each location choice. For example, with $\beta = 2, \eta = -.963, \nu = .75, \rho = -.01$, and $y = \$50,000$, the inequalities for a household living in community B reduce to:

$$\begin{aligned} .753228 + .994746 \cdot \alpha_{i0} &\leq .750448 + .997771 \cdot \alpha_{i0} \\ .753228 + .994746 \cdot \alpha_{i0} &\leq .753824 + .994318 \cdot \alpha_{i0} \quad . \\ .753228 + .994746 \cdot \alpha_{i0} &\leq .756416 + .993092 \cdot \alpha_{i0} \end{aligned} \quad (8)$$

Solving this system reveals that households in community B must have values for α_0 between 0.92 and 1.39. Any value for α_0 within this range is consistent with the choice of community B . Thus, the choice of community B “set identifies” α_0 to be within that range. Figure 2 displays the full partition of α_0 for households with income equal to \$50,000. If income were to increase, the borders of the partition would shift to the left.⁶ Intuitively, the choice to live in B reveals a lower range of preferences for a wealthier household that can more easily afford to live in communities C and D , which are more expensive because they provide more public goods.

Figure 2 also illustrates two key limitations of set identification. First, preferences are not point identified within the bounds of a set. The choice of community B reveals only that the household's preferences lie somewhere in the range: $0.92 \leq \alpha_{i_0} \leq 1.39$. Second, the preference sets that correspond to communities with extreme provision of public goods are not bounded by the revealed preference logic in (5). The preference set for community A is not bounded from below and the set for community D is not bounded from above. Therefore, a distribution must be specified for α_0 in order to point-identify preferences. This added information transforms the observed location choices by a population of households into a distribution of preferences.

Recent applications have sought to relax vertical differentiation to allow households to differ in their relative preferences for different amenities [5, 14]. This is often described as “horizontal differentiation”. In this case, observed location choices are required to set-identify more heterogeneous preference parameters. This generalization complicates the partitioning process. Returning to the CES utility function, the inequalities in (5) now implicitly define an infinite number of $(\alpha_{i_0}, \alpha_{i_1}, \alpha_{i_2})$ points that would make a household exactly indifferent between each pair of communities. Figure 3 shows the “indifference loci” formed from these points in \mathfrak{R}_+^2 , using the restriction that $\alpha_{i_1} + \alpha_{i_2} = 1$ to depict the loci in α_0, α_1 space.⁷ A household with preferences anywhere along the locus from (0.34, 0) to (4.25, 0.57) is exactly indifferent between communities B and C . Households above this locus have stronger relative preferences for g_1 , and therefore prefer community B . Households below the locus have stronger relative preferences for g_2 , and therefore prefer community C . The indifference loci delineate 12 areas, each of which corresponds to a different ranking of communities by the utility they provide. The rankings (from highest to lowest) are shown in parentheses in the figure. Notice that adjacent areas often have the same utility maximizing community. Community A maximizes utility

throughout areas 1 and 2, for example. By taking the union of areas that share the same utility-maximizing community, Figure 4 partitions preference space into four regions that rationalize each of the four community choices. For example, taking the union of the two areas that have community *A* as the top-ranked choice (areas 1 and 2 in Figure 3) produces region *A* in Figure 4. A household with preference anywhere in region *A* will maximize its utility by choosing to live in community *A*.

The partitioning process can serve to illustrate how the identifying power of the indirect utility function differs under vertical and horizontal differentiation. For example, Figure 5 overlays the vertical partition from Figure 2 as a dashed line on top of the horizontal partition from Figure 4. The resulting comparison illustrates how preference sets revealed by vertically differentiated sorting are subsets of their horizontally differentiated counterparts. The dotted vertical lines mark the intersection between the vertical preference sets and the horizontal preference sets. For example, in the vertical case the choice of community *B* indicates that the household's preferences belong to the set: $(\alpha_1 = .46, 0.92 \leq \alpha_0 \leq 1.39)$. This is exactly the portion of the dashed line in Figure 5 which intersects the lower left portion of region *B*—the preference set identified by the choice of community *B* in the horizontal case.

The process of partitioning preference space also illustrates how the econometrician's depiction of preference heterogeneity controls the scope of substitution patterns. Let community *k* be a substitute for *j* if and only if $\partial q_k / \partial p_j > 0$.⁸ With vertical differentiation each community has at most two substitutes, the adjacent communities in the ranking by public goods. With horizontal differentiation the total number of substitutes for each community falls between 2 and *J*, depending on the number of choices relative to the number of public goods [1]. The communities that are substitutes share “borders” in the partition of preference space.

Community B , for example, shares borders with each of the other three communities in Figure 4. Consider a marginal increase in the price of housing in community B . Households that currently reside in B but have preferences on the border between B & D will respond to the price increase by moving to community D . Likewise, households on the borders between B & A and B & C will move to communities A and C . Graphically, the price increase “pushes” the borders of region B inward so that it contracts while every other region of the partition expands.

The substitution patterns that arise from horizontal differentiation have intuitive properties. Locations that are similar in terms of prices and public goods are more likely to be substitutes than those that are not. Notice that in Figure 4 the two communities with intermediate levels of public goods, B and C , share borders with each of the other three locations while the most and least expensive locations, A and D , do not share a border. Because locations A and D are furthest removed in terms of prices and public goods, it seems natural to expect that there are few, if any, households that consider them as close substitutes.

Thus far, we have seen how moving from general axioms on consumer preferences to a parametric choice for the indirect utility function allows Varian’s [29] graphical approach to bounding indifference curves to be extended into preference space. While a partition of preference space provides more information about household preferences, it is still insufficient to develop welfare measures for policy changes. Translating the partition into welfare measures requires adding distributional assumptions.

3.2. *Distributional Assumptions and Point Identification of Welfare Measures*

In an empirical analysis, distributional assumptions will influence estimated welfare measures for policy changes. For example, the marginal-willingness-to-pay (MWTP) for a

public good based on the utility function in (6) can be expressed as (9).

$$MWTP(g_1) = \frac{\partial V / \partial g_1}{\partial V / \partial Y} = \frac{\alpha_{i1} \alpha_{i0} (\bar{g}_{ij})^{\rho-1} y_i^v}{V_i^\rho - \alpha_{i0} (\bar{g}_{ij})^\rho}. \quad (9)$$

Since the MWTP is a function of $(\alpha_{i0}, \alpha_{i1}, \alpha_{i2})$, every point in the partition of preference space corresponds to a specific MWTP. Suppose we want to infer the distribution of MWTP for g_1 for households living in community C . The choice of community C reveals only that households living there have preferences somewhere in region C of the partition. To calculate their distribution of MWTP we must first specify a function for the preference parameters over region C . Two extreme cases provide bounds for the MWTP. The first case is where every household has preferences at the point (#) in Figure 5, which corresponds to the lowest MWTP of any point in region C . The opposite extreme is where every household has preferences at (*), which corresponds to the highest MWTP. Thus, $[MWTP(\#), MWTP(*)]$ spans the range of possible measures for individual MWTP. The wider this range the greater the sensitivity of welfare effects to the distributional assumptions made in order to move from set to point identification.

The process of partitioning preference space also reveals there are limits to what can be learned about welfare measures from revealed preference analysis. Consider community D in Figure 5. Because it provides the most public goods, it will attract households with the strongest preferences and the highest MWTP. These households may make the largest contribution to summary measures of the average MWTP. In this case, revealed preference analysis is limited because there is no upper bound on α_0 in region D of the partition. To recover a MWTP distribution for community D , either an absolute upper bound must be imposed on α_0 or a distribution that limits weight in the tail.

Figure 5 also illustrates how the econometrician's choice for how to represent preference

heterogeneity resembles a bias/variance tradeoff with respect to welfare measures. Suppose that horizontal differentiation is the “true” form of preference heterogeneity. By restricting relative preferences, vertical differentiation biases welfare measures. Introducing horizontal differentiation eliminates the restriction that causes bias, but the added dimensionality of preferences increases the scope for distributional assumptions to influence results.

The welfare implications of distributional assumptions and the depiction of preference heterogeneity increase with the size of the change in public goods being considered. A sufficiently large shock to the provision of public goods would induce some households to move to a new home which, in turn, would require rents to adjust in order to clear the market. Seig et al. [24] develop a strategy for using the information about preferences recovered during the estimation to simulate this “general equilibrium” adjustment process.⁹ The prices and location choices that define the new locational equilibrium can then be used to generate a distribution that describes households’ *ex post* willingness-to-pay for the change. This strategy has since been extended to recognize that the provision of public goods such as open space [30] and opportunities for social interaction [5] may also be affected by the adjustment process. In general equilibrium analyses such as these, the depiction of preference heterogeneity plays two distinct roles in the determination of welfare measures. First, it interacts with distributional assumptions to determine what is learned about preferences during the estimation (as in the partial equilibrium case). Second, the substitution patterns implied by the depiction of preference heterogeneity govern the migration patterns when households relocate during the simulation. Both roles will influence welfare measures for policy changes.

3.3. *Unobserved Characteristics*

The preceding analysis assumed the econometrician observes every characteristic which influences households' market choices. This is a serious departure from empirical work where observable characteristics are often insufficient to explain observed behavior. To illustrate the problem, consider adding a fifth community, E , to the choice set in Table 1, which has $p_E = 0.9$ and $g_{1E} = g_{2E} = 2.1$. The specification for utility in (6) implies that any household with strictly positive marginal utility for the two public goods would choose to live in community E . Does this necessarily imply that households living in communities A , B , C , and D obtain negative marginal utility from public goods? Another explanation is that their housing prices reflect important characteristics which are not observed by the econometrician.

Lancaster [15] recognized that unobserved characteristics pose a problem for the characteristics approach to consumer theory. He suggested that the analyst continue adding characteristics to the model until every observed choice can be explained. While his reasoning is sound it is also conceivable that in any given application there will be some important characteristics that simply cannot be observed. For example, in their application to computers, Bajari and Benkard [2] find that even after controlling for an exhaustive set of 19 characteristics, more than half of observed choices cannot be rationalized by their linear utility function.¹⁰ A variety of econometric strategies have been developed for this situation [2, 7, 10]. To relate these strategies to the central issue of identifying the values consumers assign to public goods, we must contrast the restrictions which each econometric strategy imposes on the form of heterogeneity in preferences for *unobserved* characteristics.

Let there be $M - N$ public goods which are observed by households, but not by the econometrician. To distinguish them from observed public goods, label them as ξ_{N+1}, \dots, ξ_M . Assuming each unobserved public good has an additively separable effect in the public goods

index, the linear index from equation (6) can be rewritten as (10).

$$\bar{g}_{ij} = \alpha_{i1}g_{1j} + \dots + \alpha_{iN}g_{Nj} + \bar{\xi}_{ij}, \quad (10)$$

$$\text{where } \bar{\xi}_{ij} = \alpha_{iN+1}\xi_{N+1j} + \dots + \alpha_{iM}\xi_{Mj}.$$

Econometric strategies for dealing with unobserved characteristics can be viewed as imposing one of two restrictions on $\alpha_{iN+1}, \dots, \alpha_{iM}$. The first approach is to restrict these preferences to be vertically differentiated. If households have identical relative preferences for the unobserved public goods, the i subscript can be dropped from $\bar{\xi}_{ij}$, as in (11).

$$\bar{g}_{ij} = \alpha_{i1}g_{1j} + \dots + \alpha_{iN}g_{Nj} + \bar{\xi}_j. \quad (11)$$

In this case, the $J \times 1$ vector $\bar{\xi}$ can be treated as a structural error term and recovered during the estimation assuming $\bar{\xi}_j$ and g_j are independent, or instruments are available [2, 10]. More generally, condensing all the unobservables into a single term which varies across choices, but not consumers, has been labeled the “pure characteristics” approach to modeling the demand for a differentiated product [8].

The second econometric strategy relaxes vertical differentiation to allow households to differ in their relative preferences for unobserved characteristics. When households have horizontally differentiated preferences for ξ_{N+1}, \dots, ξ_M , their net effect on the public goods index will be unique for each (household, location) combination. In this case, the public goods index in (10) can be rewritten as (12) after introducing a slight change in notation: $\bar{\xi}_{ij} \equiv \bar{\xi}_j + \varepsilon_{ij}$.

$$\bar{g}_{ij} = \alpha_{i1}g_{1j} + \dots + \alpha_{iN}g_{Nj} + \bar{\xi}_j + \varepsilon_{ij}. \quad (12)$$

As in the pure characteristics model, $\bar{\xi}_j$ can be interpreted as the mean utility from the unobserved characteristics associated with location j . The second term, ε_{ij} , is the consumer-

specific deviation from that mean.¹¹ Notice that equation (12) looks like a random utility model (RUM). Indeed, imposing the usual *iid* extreme value assumption on ε_{ij} would make it possible to express the probability that household i chooses location j as a function of the parameters which characterize the distribution of $\alpha_{i1}, \dots, \alpha_{iN}$. This approach has formed the basis for studies which convert (12) into an indirect utility function by adding linear terms for housing prices and structural characteristics (e.g. [5]).¹²

The conceptual framework for partitioning preference space and decomposing welfare measures can be applied to both the RUM and pure characteristics models. From the perspective of identifying welfare measures, the key difference between them lies in what they imply for the dimensionality of the partition. Relative to the pure characteristics model, the RUM framework adds J dimensions to the partition. For example, if the public goods index in (6) were replaced with (12), the partition of preference space associated with the 5-community example would expand by five dimensions: $(\varepsilon_{iA}, \varepsilon_{iB}, \varepsilon_{iC}, \varepsilon_{iD}, \varepsilon_{iE})$. While this 8-dimensional partition is difficult to visualize, it will have two obvious properties. First every region will be unbounded. For example, region B is clearly unbounded in the ε_B dimension because increasing ε_{iB} always increases the utility from locating in community B . Given values for the other parameters $(\alpha_{i0}, \alpha_{i1}, \alpha_{i2}, \varepsilon_{iA}, \varepsilon_{iC}, \varepsilon_{iD}, \varepsilon_{iE})$ there must be some threshold for ε_B above which community B will always maximize utility. The second distinguishing feature of the partition is that every region must share a border with every other region. Within region A for example, ε_B can always be increased until communities A and B provide exactly the same utility, and the same is true for $\varepsilon_C, \varepsilon_D$, and ε_E . This implies every community is a substitute for every other community.

It seems intuitively plausible that households will differ in their preferences for individual

locations in ways the econometrician cannot observe. A household may be attached to a particular location because their family and friends live there for example. By failing to account for these effects, pure characteristics models may generate biased welfare measures. However, the extra flexibility provided by adding J idiosyncratic error terms comes at a considerable cost because when every region of the partition is unbounded, so is the range of welfare measures associated with every location choice. The endpoints of these ranges are only identified by the distributional assumptions on ε_{ij} which, therefore, have the potential to dominate what we learn about willingness-to-pay for environmental amenities and other non-market goods. Thus, selecting the form of heterogeneity in preferences for unobserved characteristics resembles another bias/variance tradeoff with respect to welfare measures.

If the idiosyncratic error terms in a random utility model are primarily seen as capturing the cost of moving away from family, friends, and a familiar environment, one potential compromise between the RUM and pure characteristics frameworks would be to modify the choice set in a pure characteristics model to include moving costs.

3.4. Defining the Choice Set and Relaxing Free Mobility

It is well known that the results from discrete-choice models can be sensitive to the way the econometrician defines the choice set. However, previous efforts to investigate this sensitivity have been conditioned on specific distributional assumptions.¹³ The concept of partitioning preference space provides an opportunity to generalize this line of work by analyzing the extent to which set identification of welfare measures depends on the econometrician's decisions about how to define the study region and how to subdivide it into a set of housing communities. It also provides an opportunity to investigate how the standard free

mobility assumption contributes to the identification of welfare measures.

Recall that the revealed preference logic that underlies the partitioning process relies on the assumption that households are free to move across the urban landscape. With this logic in mind, the study region should be sufficiently small that moving costs can be ignored and the region should be isolated from outside alternatives in a way that substantially increases the cost of leaving. If the study region is not isolated from outside alternatives in an obvious way, the analyst must find other criteria to define its boundaries. This process poses a tradeoff. Expanding the study region provides more information to identify preferences at the cost of greater reliance on the free mobility assumption.

To see how expanding the geographic borders of the study region provides more information about preferences, consider the partition in Figure 6.a. It corresponds to a study region based on three of the four communities in Table 1 (*A*, *C*, and *D*). Comparing this partition with the one based on all four communities (reproduced as panel b) shows that adding community *B* to the choice set diminishes the region of preference space assigned to households in *A*, *C*, and *D*. Thus, the two partitions illustrate Samuelson's [22] observation that consumer preferences are identified by the choices they make and by the choices they could have made, but did not. It follows from Figure 5 that changing the shape of the partition will also change conclusions about the MWTP of households who live in the corresponding communities.

Once the study region has been defined, it must be subdivided into a set of housing communities. The assumptions that underlie the theoretical model in section 2 provide some general guidelines for the subdivision but in practice the analyst is usually required to make decisions that can influence the resulting estimates.¹⁴ By now it should be obvious that incorrectly aggregating communities will change the shape of the partition in ways that can bias

welfare measures. This can be seen by comparing the partition in Figure 6.b. with the one in Figure 6.c. Assuming panel b represents the correct aggregation, the comparison shows the consequences of incorrectly aggregating communities *A* and *C*. The price of housing and provision of air and school quality in the new community *AC* is calculated as a simple average of their levels in communities *A* and *C*. This aggregation affects the set of values for MWTP which are consistent with the choice to live in each community.

How would relaxing free mobility change these results? If moving costs prevent households from locating in the community that would be their preferred choice, were they freely mobile, the resulting locational equilibrium will be hysteretic. Graphically, hysteresis would imply that the borders of the partition overlap. Figure 7 illustrates this using the four-community example where the annualized cost of moving is set to \$350.¹⁵ Regions of the partition denoted by the letters *A*, *B*, *C*, and *D* have the same interpretation as in earlier partitions. As before, any household with preferences in region *C* will maximize its utility by locating in community *C*. Households with preferences in the 7 regions of the partition that have labels beginning with the letter *C* (other than region *C*), such as *CB*, *CBD*, and *CDB*, would choose to live in community *C* if moving costs were zero. However, if they were born in the communities denoted by the other letters in the labels, the annualized cost of moving would prevent them from migrating to *C*. Likewise, some households living in community *C* may have preferences in regions of the partition like *BDC*, indicating that they prefer the combination of public goods and housing prices provided by communities *D* and *B* but are not willing to pay the cost of relocating there. By expanding the size of the preference sets capable of explaining each observed choice, introducing moving costs increases the sensitivity of welfare measures to distributional assumptions. In other words, moving costs increase the range of consistent

explanations for observed location choices, which decreases the precision with which we will be able to estimate the willingness-to-pay for an environmental amenity. Therefore, treating households as if they were freely mobile will create a false sense of precision in estimates for welfare measures.

Figure 7 also illustrates how data on migration patterns can provide more information to identify the preferences of households who move.¹⁶ Consider two households currently living in community *C*, one that was born there and another that migrated from *B*. For the household born in *C*, the logic is the same as before. The decision to remain in their native community implies that their preferences must lie somewhere within the contiguous area formed by the 19 regions that contain the letter *C*. The decision not to move does not help to identify their preferences. In contrast, preference sets shrink for households that move. By choosing to move out of community *B*, the second household reveals that its preferences do not lie in any region of the partition containing the letter *B*. This leaves only three regions capable of explaining their choice to move to community *C*: regions *CA*, *C*, and *CD*. The union of these three regions is smaller than the preference set assigned to households living in community *C* in the version of the model with zero moving costs. The difference is defined by the area:

$CAB \cup CBA \cup CB \cup CBD \cup CDB$. In an application where some households move and others do not, introducing moving costs may increase or decrease the precision in estimates for average willingness-to-pay. The net effect will depend on the share of households who move and on the shape of the partition.

In summary, this section has illustrated how the econometrician's choices for the parametric form of the indirect utility function, the depiction of preference heterogeneity, and the choice set all interact to jointly identify a partition of preference space which rationalizes

observed behavior. The borders of the partition reflect the substitution patterns conveyed by the form of heterogeneity. The regions of the partition provide set identification of preferences. Assumptions about the distribution of preferences within each region are what finally allow the econometrician to generate point estimates of welfare measures for policy changes. Together, these observations provide a framework for decomposing the identification of structural estimates for non-market values. The next section uses this framework to present some initial evidence on the relative economic importance of the different identifying assumptions.

4. Decomposing the Identification of Willingness-to-Pay for Improved Air Quality

The observation that the heterogeneous preference parameters are set identified without distributional assumptions suggests an empirical strategy for decomposing the identification of non-market values: first recover the region of preference space that justifies each observed location choice, then sample from that region according to alternative distributional assumptions. An upper bound distribution on welfare measures can be produced by assigning every household the highest possible willingness-to-pay that would be consistent with its observed location choice. Assigning every household their lowest possible willingness-to-pay generates a lower bound distribution. Together, these two distributions bound the distribution of welfare measures under any assumptions about the distribution of preferences. This bounding approach disentangles the relative identifying power of restrictions on the utility function and assumptions on the distribution of preferences. It can be implemented using the two-step estimators developed in [2] and [14].

A brief overview of the two-step approach to estimation can help to explain the decomposition. The estimation begins by recovering all the homogeneous parameters which

enter the utility function. For the pure characteristics version of equation (6) these include β, η, ν, ρ , and $\bar{\xi}$. The housing demand parameters can be estimated from the empirical analog to the demand function in (7), and ρ and $\bar{\xi}$ can be recovered through an iterative process that selects values for these parameters which match predicted and observed behavior. Once the homogeneous parameters have been recovered, a Gibbs sampling algorithm can be used to draw a random sample of uniformly distributed points from each region of the partition. Given a starting value in a particular region of the partition, the Gibbs algorithm will sample uniformly within that region. Let the starting value be defined by a point: $(\alpha_0^0, \alpha_1^0, \dots, \alpha_N^0)$. The algorithm iterates over the following steps:

- (0) Given $(\alpha_1^0, \alpha_2^0, \dots, \alpha_N^0)$, solve equation (5) for the minimum and maximum bounds on α_0 , and use those bounds to take a new draw, α_0^1 .
- (1) Given $(\alpha_0^1, \alpha_2^0, \dots, \alpha_N^0)$, solve equation (5) for the minimum and maximum bounds on α_1 , and use those bounds to take a new draw, α_1^1 .
- .
- .
- .
- (N) Given $(\alpha_0^1, \alpha_1^1, \dots, \alpha_{N-1}^1)$, solve (5) for the minimum and maximum bounds on α_N , and use those bounds to take a new draw, α_N^1 , resulting in $(\alpha_0^1, \alpha_1^1, \dots, \alpha_N^1)$, the first new point in the region. Using this new point, return to step (0).

The points generated by the Gibbs algorithm effectively trace out the shape of the region. When this process is repeated for each region, the resulting collection of points will approximate the partition of preference space. Finally, converting each point into the corresponding willingness-to-pay provides the information needed to generate distributions of welfare measures. The remainder of this section uses the Gibbs algorithm to decompose Kuminoff's [14] estimates of the marginal willingness-to-pay for improved air quality in Northern California.

4.1. The Northern California Equilibrium Sorting Model

The Northern California model describes households living in the region's two largest population centers: the San Francisco and Sacramento Consolidated Metropolitan Statistical Areas (CMSA). Together, these two CMSAs contain about 9 million people, roughly 25% of California's population and 3% of the U.S. population. The region is divided into 122 housing communities, each of which is distinguished by its price of housing and by its provision of two public goods, air quality and school quality. As in most applications, housing communities are defined as unified school districts.¹⁷ School quality within each district is measured using the Academic Performance Index (API), a composite index of standardized test scores reported by the California Department of Education. Air quality is measured using a spatial weighted average of ozone concentrations reported by the California Air Resources Board.¹⁸ Finally, an index of housing prices was estimated from a log-linear parameterization of the hedonic price function in (3) using micro data on the sale price and structural characteristics of 540,642 individual homes sold in the region between 1995 and 2005. The estimation was based on the horizontally differentiated utility function in (6), using the public goods index in (11).

The dimensions of the choice set and estimates for the homogeneous parameters are

representative of the literature. Previous applications have used between 91 and 150 communities, and most have included two or three amenities in the public goods index. Likewise, estimates for the price elasticity ($\hat{\eta} = -0.38$), income elasticity ($\hat{\nu} = 0.66$), and elasticity of substitution ($\hat{\rho} = -0.118$) have the typical signs and magnitudes.¹⁹ The estimated $\bar{\xi}$ vector implies that the composite effect of unobserved public goods becomes increasingly important as one moves closer to the San Francisco Bay. For a full description of the data, estimation, and results see [14]. Here, the goal is simply to use those results to demonstrate the economic importance of different maintained assumptions in an equilibrium sorting application.

Table 2 reports summary statistics for the choice set used in the partitioning process, by CMSA. While the two CMSAs are adjacent, their major business districts are 80 to 120 miles apart, with the San Francisco CMSA containing 101 of 122 communities and 79% of the region's population. Compared to the Sacramento CMSA, the San Francisco communities tend to have lower ozone concentrations, higher values for the API, higher housing prices, and higher values for median household income. There is also more variation in these characteristics within San Francisco. The distinction between the two CMSAs provides an opportunity to investigate the sensitivity of welfare measures to alternative definitions for the choice set.

4.2. Bounding Estimates of the MWTP for Improved Air Quality

The Gibbs sampling algorithm was used to recover an approximation to the partition of preference space by drawing 2000 uniformly distributed points from each region of the partition. Each point was converted into the annualized willingness-to-pay for a marginal (1 part per billion) reduction in ozone concentrations (MWTP).²⁰ Table 3 reports average MWTP at different geographic levels of aggregation under different distributional assumptions. Estimates

in the first three columns are based on the usual free mobility assumption. These estimates reveal that average MWTP must lie somewhere between \$67 and \$179 regardless of the econometrician's maintained assumption about the distribution of preferences. That is, assigning every household the lowest possible value for MWTP that would be consistent with its observed location choice results in an average MWTP of \$67. Similarly, assigning every household its highest possible MWTP implies an average of \$179. Any other assumption about the distribution of preferences will produce an estimate within this range. For example, when preferences are assumed to be uniformly distributed within each region of the partition average MWTP is \$123.

It is difficult to make a direct comparison between the range of estimates for average MWTP in table 3 and the range of estimates in the hedonic literature where air quality is typically measured by particulate matter or by the number of days that ozone concentrations exceed state or federal standards. However, to the extent that ozone, particulate matter, and the frequency of air quality violations are simply different proxies for clean air, they can be compared in terms of a common proportionate change. Using this logic to convert the range of estimates in the existing hedonic literature to measures that would be comparable to the average willingness-to-pay for a 1 ppb reduction in ozone concentrations implies a range from \$7 to \$154.²¹ The estimates in table 3 seem relatively high by comparison. This may reflect methodological differences, or simply differences in the study region. For example, the upper bound of \$154 from the hedonic literature is from Southern California where ozone concentrations are among the highest in the nation. Since ozone concentrations are generally lower in Northern California, one would expect households with stronger preferences for air

quality to locate there. The same sorting logic can also be used to explain variation in estimates for MWTP within Northern California.

Since San Francisco communities tend to have higher air quality, higher housing prices, and higher median income than their Sacramento counterparts, it makes sense that the upper and lower bounds on average MWTP in the San Francisco CMSA (\$75 to \$200) are higher than in Sacramento (\$34 to \$102). The identifying power reflected in these bounds varies across individual communities.²² For example, average MWTP is tightly bounded for the La Honda-Pescadero Unified school district (\$165 to \$176). La Honda-Pescadero has approximately average values for the price of housing and each public good. Other communities with similar characteristics limit the region of preference space that would be consistent with the choice to live there. In comparison, Sacramento City Unified has below average values for each characteristic, few close substitutes, a larger preference region, and less informative bounds on average MWTP (\$41 to \$118).²³ Figure 8 compares the Gibbs approximations to the preference regions for each community. Since La Honda-Pescadero is more expensive and provides more of each public good than Sacramento City, the choice to live there reveals stronger overall preferences for public goods (α_0), which translates into a higher MWTP.

The middle three columns of Table 3 report the results from using a different definition for the choice set to partition preference space. Specifically, San Francisco communities are excluded from the choice set faced by Sacramento households, and vice versa. This implies the cost of moving between the two CMSAs is effectively infinite. Notice there is virtually no effect on the MWTP in San Francisco. This is because San Francisco communities tend to be closer substitutes with each other than with Sacramento communities. Graphically, San Francisco communities tend to border each other in the “free mobility” partition of preference space.

Therefore, adding Sacramento to the choice set provides very little additional information to identify MWTP for improved air quality in San Francisco. The reverse is not true. Without San Francisco in the choice set, the upper bound on average MWTP in Sacramento nearly doubles. Intuitively, the choice not to live in the communities with the highest air quality, which are located in San Francisco, helps to identify an upper bound on the willingness-to-pay for improved air quality by households in Sacramento.

Finally, the last three columns of Table 3 report the results from an intermediate case where households are free to move between San Francisco and Sacramento, but must pay a \$1000 annual cost of doing so. This fee could represent higher commuting costs, a salary reduction associated with a new job, or the psychological cost of moving to a less familiar environment. The \$1000 figure is simply meant to illustrate how moving costs would affect the results. Not surprisingly, relaxing the free mobility assumption, while retaining all the information in the choice set, provides less information about MWTP than when mobility is assumed to be free, but more information than when households are assumed to be immobile.

5. Conclusions

Structural assumptions matter for benefit measurement. This paper has illustrated how the assumptions of a structural model influence estimates for partial equilibrium welfare measures. The parametric form of the utility function, the depiction of preference heterogeneity, and the definition for the choice set were shown to jointly identify a partition of preference space which explains observed behavior. The borders of the partition reflect the substitution patterns conveyed by the form of heterogeneity. The regions of the partition set-identify preferences. Maintained assumptions about the distribution of preferences within each region are what point-

identify welfare measures for policy changes. Together, these observations provide a new framework for decomposing the identification of non-market values. This framework is not limited to equilibrium sorting models of household location choice. It applies to the broad class of structural discrete-choice models used for nonmarket valuation which also includes random utility models, structural hedonic models, and corner solution models.

Using the new framework to decompose estimates of the MWTP for improved air quality in Northern California revealed that the identifying power of assumptions about the distribution of preferences can vary widely within a single market. Distributional assumptions were capable of changing estimates for average MWTP by more than 200% in Sacramento City and by less than 7% in La Honda-Pescadero. This information would be useful in a benefit-cost analysis. For example, if we were to evaluate the benefits to Northern California households from the new federal standards on ambient ozone concentrations, we could assure policymakers that our estimates for La Honda-Pescadero are not driven by seemingly arbitrary distributional assumptions. We could also use the partition of preference space to provide “best case” and “worst case” scenarios for the benefits in Sacramento City. In both cases, using the partition to report bounds on welfare measures would make their identification more transparent. This would begin to address the Office of Management and Budget’s guidelines for quantifying uncertainty in the modeling assumptions which enter benefit-cost analyses of federal programs.

Of all the maintained assumptions which underlie discrete-choice analysis, the statistical distributions used to characterize unobserved preference heterogeneity seem the most difficult to justify on the basis of economic theory, market data, or common sense. Therefore, using the partition of preference space to analyze the sensitivity of welfare measures to distributional assumptions would be a natural first step toward greater transparency in modeling the demand

for environmental amenities. One of the next steps would be to relax parametric restrictions on the shape of the utility function. For example, in some applications it may make sense to interact public and private goods so that utility is not separable in either component.

Finally, the partition of preference space was used to illustrate how failing to account for the cost of moving between different geographic areas introduces a false sense of precision into estimates for welfare measures. This bias can be eliminated if data on moving costs and migration patterns are available. For working households, moving costs are directly connected to the choice of a job. Moving to a new home in a higher amenity area may require taking a job with a lower wage or a less desirable commute. Extending the equilibrium sorting literature to model the tradeoffs faced by working households making a joint job-house choice would be an important topic for further research. Of course, free mobility is not strictly a structural assumption. It is also one of the assumptions which allow us to interpret partial derivatives of the hedonic price function as measures of the marginal willingness-to-pay. Thus, another topic for future research would be to investigate how moving costs and sorting behavior influence what we stand to learn about welfare measures from reduced form hedonic models.

Endnotes

¹ OMB suggests testing the sensitivity to alternative modeling assumptions and then reporting probability distributions of benefits, costs, and net benefits. This discussion can be found under “Treatment of Uncertainty”.

² Intuitively, freely mobile households will sort themselves across the urban landscape until any two locations with identical housing characteristics and public goods also have identical prices.

³ Given a parametric form for (3) and data on housing transaction prices and their structural characteristics, the price of housing in each community can be recovered empirically as a community-specific fixed effect [23].

⁴ The notion that a set of differentiated communities can be unanimously ranked by quality is analogous to Lancaster’s [16] description of “vertical” product differentiation. In contrast, when households differ in their relative preferences for public goods, they may also differ in the way they rank communities—a situation analogous to Lancaster’s description of “horizontal” product differentiation.

⁵ Ekeland et al. [9] outline restrictions on utility functions and the data generating process that allow them to identify preferences from the hedonic price function in a single market. One restriction is that choices are differentiated by a single non-price characteristic. Vertical differentiation provides a way to interpret this dimensionality restriction by ensuring that all the non-price heterogeneity in choices enters utility through a single choice-specific variable, \bar{g}_j .

⁶ Alternatively, the borders of the partition will be invariant to income if it enters utility through an additively separable term. In this case, income drops out of the revealed preference inequalities in (5).

⁷ Each locus is drawn from 100 $(\alpha_0, \alpha_1, \alpha_2)$ points. Unlike the vertically differentiated case shown in (8) these points cannot be expressed analytically. Each point was identified by solving a non-linear rootfinding problem in α_0 , given values for α_1 and α_2 . These values were drawn from a grid with 100 equally spaced points between (0,1) and (1,0) in (α_1, α_2) space, using the restriction that $\alpha_{i1} + \alpha_{i2} = 1$.

⁸ This corresponds to Anderson et al.'s [1] definition of “strong gross substitutes” in a discrete-choice model.

⁹ As part of this process, the analyst must estimate or calibrate the supply flexibility of housing.

¹⁰ While the share of unexplained choices has not been reported in previous equilibrium sorting models of household location choice, the number of “rank violations” reported in Table 6 of [23] implies a lower bound of 25%.

¹¹ Decomposing $\bar{\xi}_{ij}$ into a choice-specific mean and a consumer-specific deviation from that mean provides a convenient setup for addressing endogenously determined unobserved characteristics (e.g. [5, 7]). In discrete-choice applications where endogeneity is not a major concern, it is more common to avoid the decomposition by simply defining $\bar{\xi}_{ij} \equiv \varepsilon_{ij}$. For examples, see [18] or [20].

¹² See Bayer and Timmins [6] for a discussion of IV estimation strategies for random-utility models of location choice with endogenously determined unobserved characteristics.

¹³ For an example and citations to the literature, see [18].

¹⁴ Recall that provision of public goods is assumed to be homogeneous within each community, while housing characteristics vary continuously. These two assumptions create tension in the

optimal size for communities because the first assumption is easier to satisfy with smaller communities, and the later by larger communities.

¹⁵ Deriving this partition requires subtracting \$350 from income on the right-hand side of (5).

This approach assumes an identical cost of moving from the household's current location to any other location. Alternatively, if moving costs are intended to represent the psychological cost of moving away from family and friends, they could be defined in terms of distance from the household's hometown, as in [4].

¹⁶ I thank an anonymous reviewer for suggesting this line of reasoning.

¹⁷ Exceptions are made for primary and secondary districts that do not belong to a unified district, and for the city of San Francisco which was divided into 11 supervisorial districts. All public schools in San Francisco are incorporated into a single school district, which comprises 10% of the total population in the study area.

¹⁸ It is an average of the top 30 1-hour daily maximum reading based on data from 210 monitoring stations.

¹⁹ The income elasticity and CES parameter fall near the middle of the range of estimates reported by studies that have estimated a version of the CES utility function in (6). While the estimated income elasticity falls slightly below the range reported in previous studies (0.73 to 0.94) their 95% confidence intervals overlap.

²⁰ Recall from section 3.2 that outer regions of the partition will be unbounded. These regions were bounded during the partitioning process by imposing an upper bound of \$500 (year 2000 dollars) on individual MWTP. This figure is 7.5 times the average MWTP reported by Seig et al. [24] for the same measure of air quality in Southern California. Other plausible choices for an

upper bound are unlikely to have a significant effect on the aggregate results because only 6% of households live in communities that correspond to outer regions of the partition.

²¹ See Seig et al [24] for a discussion of the hedonic studies that define this range. The estimates discussed in their study have been converted to year 2000 dollars to make them comparable to table 3.

²² Comparing the difference between the upper and lower bounds for each household in the study region reveals that MWTP is identified to within \$25 for 16% of households, to within \$50 for 38%, and to within \$100 for 68%.

²³ In La Honda-Pescadero Unified: $price = 2.59$, $API = 686$, $ozone = .05$, $\bar{\xi} = .48$, and in Sacramento City Unified: $price = 1.20$, $API = 643$, $ozone = .054$, $\bar{\xi} = .08$.

References

- [1] S.P. Anderson, A. de Palma, and J.-F. Thisse, *Discrete Choice Theory of Product Differentiation*, The MIT Press, Cambridge and London, 1992.
- [2] P. Bajari, and C.L. Benkard, Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach, *J. Polit. Economy* 113 (2005) 1239-1276.
- [3] P. Bajari, and M.E. Kahn, Estimating Housing Demand with an Application to Explaining Racial Segregation in Cities, *J. Bus. Econ. Statist.* 23 (2005) 20-33.
- [4] P. Bayer, N. Keohane, and C. Timmins, Migration and Hedonic Valuation: The Case of Air Quality, NBER Working Paper, 2006.
- [5] P. Bayer, R. McMillan, and K. Reuben, An Equilibrium Model of Sorting in an Urban Housing Market, NBER Working Paper, 2004.
- [6] P. Bayer, and C. Timmins, Estimating Equilibrium Models of Sorting across Locations, *Econ. J.* 117 (2007) 353-374.
- [7] S. Berry, J. Levinsohn, and A. Pakes, Automobile Prices in Market Equilibrium, *Econometrica* 63 (1995) 841-890.
- [8] S. Berry, and A. Pakes, The Pure Characteristics Demand Model, *Int. Econ. Rev.* 48 (2007) 1193-1225.
- [9] I. Ekeland, J.J. Heckman, and L. Nesheim, Identification and Estimation of Hedonic Models, *J. Polit. Economy* 112 (2004) S60-S109.
- [10] D. Epple, and H. Sieg, Estimating Equilibrium Models of Local Jurisdiction, *J. Polit. Economy* 107 (1999) 645-681.
- [11] M.M. Ferreyra, Estimating the Effects of Private School Vouchers in Multi-District

- Economies, *Amer. Econ. Rev.* 97 (2007) 789-817.
- [12] J.J. Heckman, Causal Parameters and Policy Analysis in Economics: A Twentieth Century Retrospective, *Quart. J. Econ.* 115 (2000) 45-97.
- [13] M.P. Keane, Structural vs. Atheoretic Approaches to Econometrics. *J. Econometrics* (Forthcoming).
- [14] N.V. Kuminoff, Recovering Preferences for Public Goods from a Dual-Market Locational Equilibrium, Ph.D. Dissertation, North Carolina State University (Economics), 2006.
- [15] K.J. Lancaster, A New Approach to Consumer Theory, *J. Polit. Economy* 74 (1966) 132-157.
- [16] K.J. Lancaster, *Variety, Equity, and Efficiency*, Columbia University Press, New York, 1979.
- [17] D. McFadden, Conditional Logit Analysis of Qualitative Choice Behavior. in: P. Zarembka, (Ed.), *Frontiers in Econometrics*, Academic Press, New York, 1974.
- [18] G.R. Parsons, A.J. Plantinga, and K.J. Boyle, Narrow Choice Sets in a Random Utility Model of Recreation Demand, *Land Econ.* 76 (2000) 86-99.
- [19] R.B. Palmquist, Estimating the Demand for the Characteristics of Housing, *Rev. Econ. Statist.* 66 (1984) 394-404.
- [20] D.J. Phaneuf, C.L. Kling, and J.A. Herriges, Estimation and Welfare Calculations in a Generalized Corner Solution Model with an Application to Recreation Demand, *Rev. Econ. Statist.* 82 (2000) 93-92.
- [21] S. Rosen, Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition, *J. Polit. Economy* 82 (1974) 34-55.

- [22] P.A. Samuelson, Consumption Theory in Terms of Revealed Preferences, *Economica* 15 (1948) 243-253.
- [23] H. Sieg, V.K. Smith, H.S. Banzhaf, and R. Walsh, Interjurisdictional Housing Prices in Location Equilibrium, *J. Urban Econ.* 52 (2002) 131-153.
- [24] H. Sieg, V.K. Smith, H.S. Banzhaf, and R. Walsh, Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods, *Int. Econ. Rev.* 45 (2004) 1047-1077.
- [25] V.K. Smith, H. Sieg, H.S. Banzhaf, and R. Walsh, General Equilibrium Benefits for Environmental Improvements: Projected Ozone Reductions under EPA's Prospective Analysis for the Los Angeles Air Basin, *J. Environ. Econ. Manage.* 47 (2004) 559-584.
- [26] C. Timmins, If You Cannot Take the Heat, Get Out of the Cerrado...Recovering the Equilibrium Amenity Cost of Nonmarginal Climate Change in Brazil, *J. Reg. Sci.* 47 (2007) 1-25.
- [27] K.E. Train, *Discrete Choice Methods with Simulation*, Cambridge University Press, Cambridge, 2003.
- [28] United States Office of Management and Budget, *Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs*, Circular A-94, 1992.
- [29] H.R. Varian, The Nonparametric Approach to Demand Analysis, *Econometrica* 50 (1982) 945-974.
- [30] R.L. Walsh, Endogenous Open Space Amenities in a Locational Equilibrium. *J. Urban Econ.* 61 (2007) 319-344.
- [31] J. Wu, and S.-H. Cho, Estimating Households' Preferences for Environmental Amenities Using Equilibrium Models of Local Jurisdictions, *Scot. J. Polit. Economy* 50 (2003) 189-

206.

Table 1

Housing Prices and Public Goods Provision in Four Hypothetical Communities

Community	Public Goods ^a		Price of Housing
	g ₁	g ₂	
A	1.25	1.25	1.04
B	1.85	1.56	1.25
C	1.66	1.86	1.30
D	2.00	2.00	1.54

^a Higher values indicate higher quality.

Table 2

Summary Statistics for Communities in the Northern California Choice Set, by Metropolitan Area

Metro Area	Attribute	Mean	St. Dev.	Min	Max
Sacramento (21 communities)	Ozone (parts per million)	0.091	0.009	0.069	0.106
	Academic Performance Index	685	78	541	846
	Index of Unobserved Public Goods	0.12	0.09	0.00	0.35
	Index of Housing Prices	1.33	0.27	1.00	2.08
	Median Household Income	59,709	9,462	41,674	73,583
San Francisco (101 communities)	Ozone (parts per million)	0.064	0.012	0.031	0.090
	Academic Performance Index	711	96	528	941
	Index of Unobserved Public Goods	0.51	0.22	0.11	1.07
	Index of Housing Prices	2.81	1.01	1.28	6.51
	Median Household Income	80,566	18,226	45,054	130,190

Table 3

Sensitivity of Average MWTP to Assumptions about Mobility and the Distribution of Preferences ^a

Region	Free Mobility			No Mobility			Limited Mobility		
	min MWTP	uniform pref.	max MWTP	min MWTP	uniform pref.	max MWTP	min MWTP	uniform pref.	max MWTP
<i>Northern California</i>	67	123	179	65	131	200	65	124	184
<u>San Francisco Metro Area</u>	75	138	200	74	137	200	74	137	200
La Honda-Pescadero School District	165	170	176	165	170	176	165	170	176
<u>Sacramento Metro Area</u>	34	66	102	34	109	198	33	74	123
Sacramento City School District	41	81	118	45	143	240	42	97	152
Cost of moving between Metro Areas (\$)		0			∞			1000	

^a MWTP is an annualized measure. “min MWTP” assigns every household the lowest value that would be consistent with their observed location choice while “max MWTP” assigns them the highest. “uniform pref.” is an intermediate case where preferences are uniformly distributed within each region of the partition of preference space.




Fig. 1. Revealed Preference Logic in Characteristics Space




Fig. 2. Partition of Preference Space: Vertical Differentiation ($y=\$50,000$)

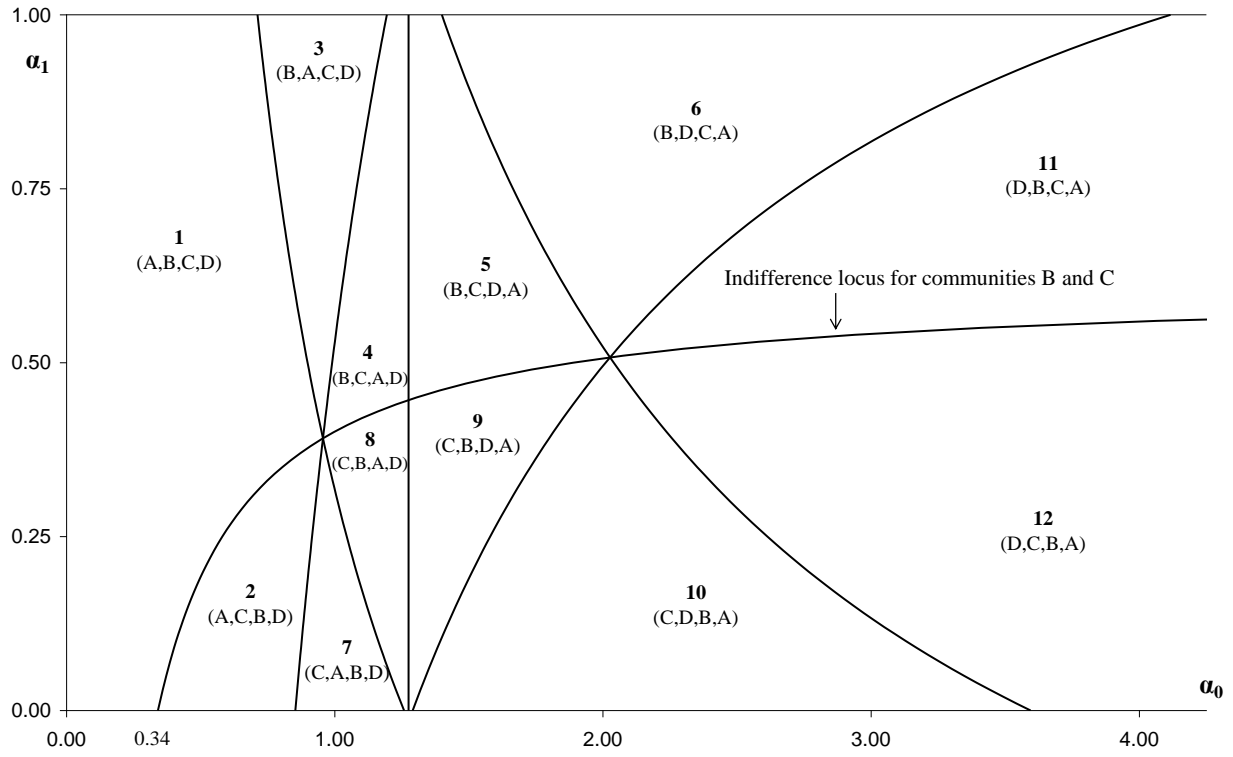


Fig. 3. Boundary Indifference Loci and Community Rankings by Utility ($y=\$50,000$)




Fig. 4. Partitioning Preference Space: Horizontal Differentiation ($y=\$50,000$)






Fig. 5. Partitioning Preference Space: Welfare Measures ($y=\$50,000$)

a. Choice Set = {A, C, D}



b. Choice Set = {A, B, C, D}



c. Choice Set = {AC, B, D}




Fig. 6. Sensitivity of the Partition to the Definition for the Choice Set ($y=\$50,000$)




Fig. 7. Partitioning Preference Space: \$350 Moving Cost ($y=\$50,000$)

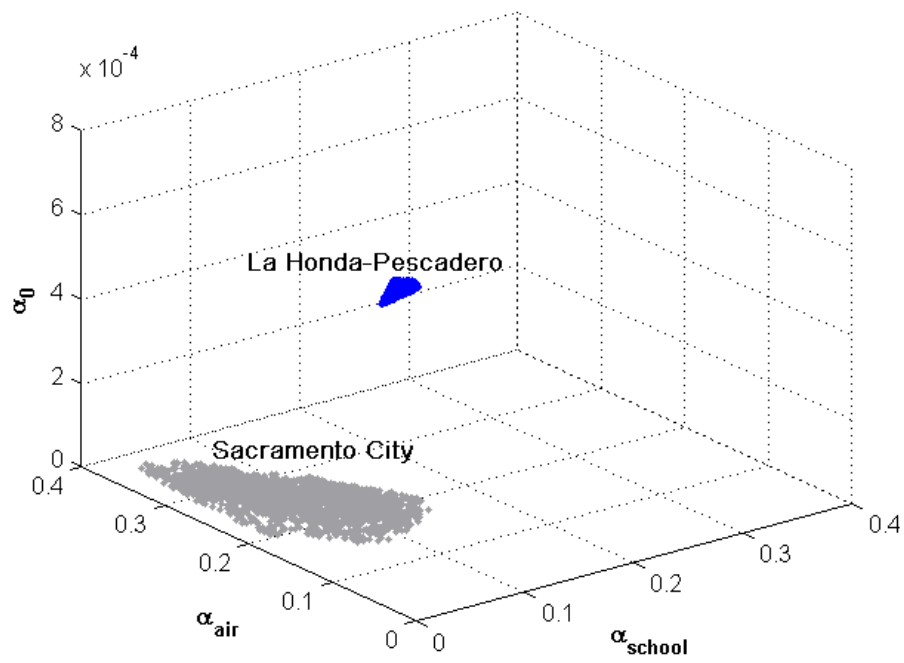


Fig. 8. Two Regions in the Northern California Partition of Preference Space ($y = \$67,500$)