Partial Identification of Amenity Demand Functions

By CONGWEN ZHANG, KEVIN J. BOYLE, AND NICOLAI V. KUMINOFF*

This paper presents a new hedonic framework for reduced form estimation of the demand for spatially delineated nonmarket amenities. We begin from a conventional model of market equilibrium where an amenity is conveyed to homeowners by virtue of their residential location choices. Different housing markets may have different hedonic price functions due to variation across markets in the joint distribution of preferences, income, regulations, and technology. In this setting, taste-based sorting within and across markets confounds point identification of reduced form descriptions for amenity demand curves. However, we demonstrate that basic knowledge of the sorting process is sufficient to construct instruments that identify bounds on demand curves. Bounds on demand curves can be translated into ranges of welfare measures for non-marginal changes in amenities. We find these ranges to be potentially informative in a demonstrative application to evaluating the benefits of improved lake water clarity in the Northeast.

January, 2015

Forthcoming in the *Journal of Environmental Economics and Management*

* Corresponding author: Kuminoff: Arizona State University, Dept. of Economics, Tempe, AZ 85287. Phone: 480.727.9802. Fax: 480.965.0748. e-mail: kuminoff@asu.edu. Zhang: Virginia Tech, Dept. of Agricultural and Applied Economics, Blacksburg, VA 24061 (e-mail: wen2007@vt.edu). Boyle: Virginia Tech, Dept. of Agricultural and Applied Economics, Blacksburg, VA 24061 (e-mail: kjboyle@vt.edu). We especially thank V. Kerry Smith for numerous suggestions on earlier drafts of this paper. We also thank the editor (Dan Phaneuf), an anonymous referee, Amy Binner, Kelly Bishop, Keith Evans, Michael Hanemann, Alex James, Steve Newbold, Kevin Novan, George Parsons, Jaren Pope, and participants of the Heartland environmental economics conference and the EPA conference on Chesapeake Bay water quality.

"Through the early history of applied hedonic analyses of urban housing markets the primary estimation problems have been (i) the interpretation of the estimated coefficients in terms of demand and supply functions, and (ii) coefficient bias introduced by omitted variables."

—Peter Linneman (1980)

1. Introduction

How little has changed! Omitted variable bias and the "second stage" identification of demand parameters are still perceived to be the two main problems with using hedonic models to estimate the demand for amenities. Both problems must be addressed in order to develop a credible estimate of the willingness to pay for a *non-marginal* change in a spatially delineated amenity. Yet, the recent empirical literature has devoted far more attention to the omitted variable problem. Numerous studies have developed ways to mitigate the bias from correlation between observed and unobserved amenities (e.g. [1,15,26,32]). In comparison, there has been relatively little progress on the identification of demand parameters. The purpose of this paper is to suggest a new approach to the main endogeneity problem with "second stage" demand estimation.

Rosen [38] suggested a simple two-stage reduced form approach to estimating demand curves for individual attributes of a differentiated good, such as housing. The first stage involves regressing product prices on product attributes (i.e. estimating a hedonic price function). The resulting estimates for marginal effects then become dependent variables in a second-stage regression of implicit marginal prices on attribute levels and demand shifters (i.e. estimating an inverse demand curve). While Rosen's first-stage model of hedonic pricing is now among the foremost tools of nonmarket valuation, his vision for a second-stage model of demand remains unfulfilled. The main problem is that taste-based sorting within and between markets for differentiated goods undermines the point identification of demand curves [4,13].

Researchers have sought to mitigate the bias from taste-based sorting by adding more information to the problem. One strategy is to assume a parametric form for the utility function (e.g. [2,3,9,19,23,39,41]). Another strategy is to collect data from multiple markets and estimate demand curves under the maintained assumption that consumers do not sort themselves across markets (e.g. [4,25,35]). A third strategy, developed by Bishop and Timmins [6] combines and extends the first two ideas.[1] While all of these proposals are intriguing, none have been widely adopted.

In this paper, we propose a fundamentally different way to extract information about the demand for a spatially delineated amenity, while returning to a simple reduced-from approach to estimation that can be implemented using data that are widely available. Our main observation is that the endogeneity problem that arises from taste-based sorting is only fatal if we limit ourselves to the extremes of point identification. By taking a broader perspective on the nature of identification, consistent with Leamer [28], Manski [30], Kuminoff [23], and Nevo and Rosen [34], we demonstrate that it is possible to identify *bounds* on demand elasticities and *ranges* of welfare measures in the presence of taste-based sorting within and between markets. Specifically, we develop a partial identification strategy for using hedonic price functions to identify consistent bounds on demand curves

[1] Bishop and Timmins track individuals who migrate across markets, write down a parametric specification for their indirect utility functions, and identify parameters of that function under the assumption that it is stable over time.

for nonmarket amenities and consistent ranges of estimates for the Marshallian consumer surplus. Our microeconometric strategy exploits Nevo and Rosen's [34] recent results on the partial identification of linear instrumental variable models.

We begin in section 2 by reviewing the endogeneity problem with hedonic demand estimation and explaining the economic intuition behind our econometric approach to the problem. Section 3 adapts Nevo and Rosen's econometric results to illustrate how basic assumptions about spatial sorting behavior are sufficient to identify bounds on demand curves and welfare measures. Section 4 provides an empirical demonstration of the main ideas. We use data on the characteristics of homebuyers and their lakefront houses in Maine, New Hampshire, and Vermont to estimate bounds on the willingness to pay for a non-marginal improvement in lake water quality. While taste-based sorting leads to uncertainty about the exact level of benefits, we find that our bounds on willingness-to-pay have the potential to be informative for policymakers. Finally, section 5 provides concluding remarks.

2. Economic Intuition for the Identification Strategy

2.1. The endogeneity problem with hedonic demand estimation

Consider a market, *m*, where households obtain utility from housing services and a composite private good, *b*. A market could be a town, city, or metropolitan area. Any house in the market may be decomposed into a bundle of private characteristics describing its structure (e.g. bedrooms, bathrooms, size of living area) and public characteristics describing its location (e.g. air quality, school quality, access to open space). We will use *x* to denote a vector of all the public and private characteristics other than the amenity of interest, which will be denoted by g. Under the standard assumptions of Rosen's [38] model, interactions between buyers and sellers define an equilibrium hedonic price schedule for the market, $p^m(x, g, \beta)$, where β is a parameter vector describing the shape of the price function.

Each individual household with income y and preferences α maximizes utility over x, g and a numeraire good, b, subject to a budget constraint defined, in part, by the hedonic price function,

$$\max_{x,g,b} U(x, g, b; \alpha) \;\; s.t. \;\; y = p^m(x, g, \beta) + b, \tag{1}$$

which produces the first order conditions from Rosen's model:

$$p_g^m \equiv \frac{\partial p^m(x,g,\beta)}{\partial g} = \frac{\partial U(x,g,b;\alpha)/\partial g}{\partial U(x,g,b;\alpha)/\partial b} \tag{2.a}$$

$$p_x^m \equiv \frac{\partial p^m(x,g,\beta)}{\partial x} = \frac{\partial U(x,g,b;\alpha)/\partial x}{\partial U(x,g,b;\alpha)/\partial b}\,. \tag{2.b}$$

The household chooses levels of each characteristic such that their willingness to pay for an additional unit (MWTP) equals its marginal implicit price. If we further assume that the marginal utility of income is constant, then the MWTP function to the right of the equality in (2a) defines the household's demand for g.[2]

$$\frac{\partial U(x,g,b;\alpha)/\partial g}{\partial U(x,g,b;\alpha)/\partial b} = D(g; x, \alpha). \tag{3}$$

Let o and u denote vectors of observed and unobserved attributes of households that underlie heterogeneity in preferences within the population. Age of the household head, income, and race are examples of demographic attributes that might be observed. Equation

[2] If the marginal utility of income is not constant then the marginal rate of substitution functions do not have the same properties as traditional Hicksian or Marshallian demand curves because the hedonic budget constraint is generally non-linear. Rosen noted this in his original paper (p.49) and McConnell and Phipps [31] provide a detailed explanation. Studies that do not explicitly assume constant marginal utility of income often refer to the marginal rate of substitution function as a "MWTP function" rather than a demand curve.

(4) uses this decomposition to define an econometric approximation to the market demand function for the amenity,

$$\frac{\partial p^m(x,g,\beta)}{\partial g} = D(g; x, o, \delta) + v(u), \tag{4}$$

where δ is a parameter vector describing the shape of the function. The econometric error term, v, depends on unobserved sources of heterogeneity in preferences.[3]

At first consideration, demand estimation seems like a simple two-step process. First, one would regress housing prices on housing characteristics to estimate β. The resulting estimates for implicit prices, $\partial p^m(x, g, \hat{\beta})/\partial g$, would then be regressed on the characteristics of houses and households in the second stage regression of equation (4) to estimate δ. The problem with this logic is that, unfortunately, the marginal price function only intersects each household's demand curve at a single point.[4] This creates an endogeneity problem for OLS estimators if households are stratified within the market according to unobserved features of their tastes.

Intuition for the resulting bias can be seen from figure 1, which is essentially reproduced from Bartik [4]. Panel A shows demand curves, D_1 and D_2, for two observationally equivalent households that make identical choices for x. However, household 2 has a stronger unobserved taste for the amenity, $v(u_2) > v(u_1)$. This induces the household to select a location within market #1 that enables it to consume more of the amenity: $g_2^* > g_1^*$. The analyst observes two market transactions, at points **a** and **b**, defining the equilibri-

[3] To focus attention on the intuition for our identification strategy, we are temporarily abstracting from other sources of error such as measurement error, functional form misspecification, and omitted variables.
[4] Brown and Rosen [8] raise a separate issue that has received considerable attention in the literature. They note that the identification of δ relies on the *seemingly* arbitrary assumption that the marginal price function is nonlinear. Fortunately, the assumption of nonlinearity is not arbitrary. Ekeland, Heckman, and Nesheim [12] clarify that nonlinearity of the marginal price function is, in fact, a generic property of hedonic equilibria.

um prices and quantities selected by each household. A line connecting these points clearly understates the slopes of the demand functions for both individuals.[5]

2.2. Market indicators as imperfect instruments

Developing instruments to resolve the endogeneity problem has proven to be difficult. Early studies addressed the problem by pooling data from multiple spatial markets and using indicators for each area as instruments (e.g. [4,7,35,43,44]).[6] In general, exogenous spatial variation in environmental amenities or development restrictions would cause equilibrium price functions to vary across spatial markets.[7] If identical households were randomly assigned to different markets, then the choices they would make when faced with different price schedules would trace out multiple points on their shared demand curve for g. Figure 1.B illustrates the intuition. The analyst now observes a second household with demand curve D_1 consuming at point **c** in market #2. A line connecting points **a** and **c** identifies D_1. This multi-market identification strategy relies on two assumptions. First, the marginal price function for g must vary across markets. Second, observationally equivalent households must *not* be stratified across markets according to unobserved features of their tastes. Intuition and empirical evidence tends to support the first assumption and contradict the second.

Since the equilibrium price function is defined by interactions between heterogeneous buyers and sellers, it is natural to expect the shapes of price functions to vary across mar-

[5] See Epple [13] for a more general characterization of the problem in a simultaneous equations framework.
[6] For example, Witte, Sumka, and Erekson [43] pool data from four cities in North Carolina: Greenville, Kinston, Lexington, and Statesville. Palmquist [35] collects data from Atlanta, Denver, Houston, Louisville, Miami, Oklahoma City, and Seattle. Zabel and Kiel [44] pool data from Chicago, Denver, Philadelphia, and Washington D.C.
[7] This follows from the observation that the hedonic price function is endogenously determined by the joint distribution of preferences, technology, and institutions inherent in a given market. As these factors vary over space or time, so does the shape of the price function [25,26].

kets due to heterogeneity in the local populations of buyers and sellers.[8] Empirical studies support this intuition. For example, Witte, Sumka, and Erekson [43] found that price functions varied across four cities in North Carolina; Palmquist [35] documented variation across eight U.S. cities; and Zabel and Kiel [44] documented variation across the four metro areas. Recent studies that have used modern econometric techniques to mitigate omitted variable bias have found similar evidence on spatial variation in the shapes of price functions estimated for different metro areas (e.g. [15,26]).

While the literature confirms that housing price functions vary across markets, related work on understanding empirical migration flows has revealed two serious problems with the "no stratification" assumption. The first problem is that long-term migration appears to be driven, in part, by preferences for amenities, especially for retired households [10,17,33].[9] The second problem is that an individual's preferences for amenities may be endogenous to their parents' location decisions. People who spend their childhood near the beach, for example, may invest in learning to surf; people who grow up near the mountains may learn to ski. These choices, and the experiences they create, may shape the individuals' preferences for coastal versus mountainous areas, which can affect their location choices later in life. Krupka [22] formalizes this logic within the broader literature on mi-

[8] Suppose there is a minimum cost of moving between two metro areas, $mc > 0$, and that the technology for producing an amenity varies across the areas due to differences in their climate, geography, or regulatory burden. Without loss of generality, denote the equilibrium price functions in the two areas by $p^1[x, g, F^1(\alpha, y, \theta)]$ and $p^2[x, g, F^2(\alpha, y, \theta)]$, where $F^m(\alpha, y, \theta)$ represents the joint distribution of preferences, income, and technology (θ) in area *m*, and $p^1[x, g, \cdot] \neq p^2[x, g, \cdot]$. Assuming utility is monotonic in the numeraire, no household in area #1 can increase its utility by moving to area #2 as long as $p^1[x, g, \cdot] - p^2[x, g, \cdot] < mc$ for all x, g. The reverse is also true. Thus, moving costs enable there to be a wedge of varying size between the equilibrium prices of identical properties located in the two areas.

[9] The potential importance of these findings is underscored by the extent of interregional migration. For example, between 1.4% and 3.2% of the U.S. population moved across state lines *every year* from 1990 through 2010. The number of migrants who crossed state lines between 2000 and 2010 was equivalent to 23% of the U.S. population, despite the sharp decline in migration after the housing market bubbles began to burst. These and other interesting statistics on migration patterns can be found at: http://www.census.gov/hhes/migration/data/cps/historical.html.

gration and human capital accumulation. He provides preliminary evidence that migrants tend to move to markets with amenities that are similar to the amenities of the markets where they were born.

Thus, if migration is influenced by preferences for amenities, or if preferences for amenities are endogenously determined through one's location-specific experiences, then observationally equivalent households will likely be stratified across markets according to their unobserved tastes. In this case, the multi-market approach to demand estimation will be inconsistent. Fortunately, one can still identify bounds on demand curves under much weaker assumptions about the sorting process.

2.2. Intuition for partial identification of demand curves for nonmarket amenities

Figure 2 illustrates the intuition for our approach to partial identification. Consider Panel A. Demand curves D_1 and D_3 describe two observationally equivalent households who are positively stratified across markets according to their unobserved tastes. Stratification is "positive" in the sense that their respective location choices generate a consumption pattern that matches their preference ordering: $v(u_1) > v(u_3)$ and $g_1^* > g_3^*$. Because the households are sorted across markets, the multi-market, instrumental-variable (IV) approach to demand estimation, discussed above, produces an inconsistent estimator of the demand curve. Graphically, a line connecting points **d** and **a** understates the slopes of D_1 and D_3. However, notice that D_1 is bounded (from above) by the IV estimator and (from below) by a perfectly inelastic demand curve through **a**. This is the economic intuition for partial identification in the presence of taste-based sorting. If we are willing to assume that households are positively stratified, then we can use market indicators to derive a consistent estimator for bounds on the demand curve. The bounds are shown by the shaded region in figure 2.A.

Figure 2.B illustrates the opposite case of negative stratification. In this case, the household with stronger unobserved tastes locates in the market where they consume less of the amenity: $v(u_3) > v(u_1)$ and $g_1^* > g_3^*$. As a result, the demand curve is bounded by the shaded region between the OLS and IV estimators. Thus, the additional information that would allow us to identify bounds on demand curves in any given application is simply an assumption about whether households are more likely to be positively or negatively stratified across markets according to their unobserved tastes for the amenity.

While positive stratification is more consistent with the standard intuition for spatial sorting behavior and may make sense for most applications, there are at least three mechanisms that could lead to negative stratification. One possibility is that preferences for different amenities are positively correlated (e.g. people like both air quality and public school quality) but the amenities themselves are negatively correlated across spatially delineated markets (e.g. one market has higher air quality and lower school quality). If households are positively stratified on one amenity (e.g. school quality), then they may be negatively stratified on the other (e.g. air quality). Negative stratification could also result from an information campaign. For example, if smog alerts in Los Angeles create an information shock that induces people to learn about the health effects of prolonged exposure to air pollution, then people living in areas with more severe air pollution may develop stronger preferences for air quality. Finally, negative stratification could arise from a strong feedback mechanism. For example, if people with strong preferences for access to open space move to areas near privately-owned open space, then the increased demand for housing may raise land values enough to stimulate the sale of privately owned open space for new residential development that reduces access to open space for the early movers, similar to the model in Walsh [42].

Importantly, the inconsistency of the conventional IV estimator increases with the strength of the between-market stratification in preferences. If we push D_1 and D_3 closer together in figures 2A and 2B, then the inconsistency diminishes. In the special case where there is no between-market stratification in preferences for the amenity of interest the conventional IV estimator point identifies demand. This could occur, for example, if the amenity of interest is both relatively unimportant in location decisions and uncorrelated with other more important drivers of sorting behavior.

3. Partial Identification of Amenity Demand Functions and Welfare Measures

In the hedonic literature, it is common to specify reduced-form approximations to market demand curves as being linear in parameters.[10] For this broad class of models, we can adapt Nevo and Rosen's [34] "imperfect instruments" framework to formalize our intuition for the econometric identification of amenity demand curves. To begin, we define an empirical counterpart to the demand function in (4) as:

$$p_g = g\delta + w\gamma + u, \tag{5}$$

where p_g denotes the implicit marginal price of the amenity and $w = [x, o]$ is a vector containing all other observed attributes of consumers and their choices. For the reasons discussed above, we would expect g to be correlated with u.

[10] For examples see [4,6,13,35,43,44].

Candidate instruments for g will be denoted by z_j $(j = 1, \ldots, J)$. The key results on identification from Nevo and Rosen [34] [henceforth NR] rely on three basic assumptions about the direction and magnitude of the correlations between g, z, and u. Following their notation we use σ_{ab} and ρ_{ab} to denote the covariance between variables a and b and the correlation coefficient respectively; likewise we use σ_a to denote the standard deviation of a.

Assumption A1: Observations on $\{p_g, g, w, z_j\}$ are stationary and weakly dependent.

Assumption A2: $E[\acute{w}u] = 0$.

Assumption A3: $\rho_{gu}\rho_{z_j u} \geq 0,\ for\ j = 1, \ldots, J.$

In general, assumption A1 ensures that the data are consistently aggregated to define a market. Stationarity requires the equilibrium price function to be stable over the dimensions of the market. Weakly dependence requires that the correlation between variables decreases as the temporal (or spatial) distance between them increases. This condition helps to avoid redundancy as the sample size grows. Overall, A1 is consistent with the existing empirical frameworks for hedonic demand estimation.[11]

Assumption A2 is made for simplicity, not necessity. NR's main results generalize to settings with multiple endogenous regressors.[12] Treating w as exogenous allows us to simplify exposition and focus our attention on the amenity of interest.

Assumption A3 is the most important. It considerably weakens the "no stratification"

[11] Kuminoff and Pope [26] suggest that stationarity of the equilibrium price function is more likely to be violated when data are pooled over long intervals (e.g. 10-20 years) spanning changes in environmental quality. With this in mind, it would be ideal to use data from a short time period when there is little scope for market fundamentals to change.

[12] All of the main econometric results go through in settings with multiple endogenous variables. The key requirements are that one has an imperfect instrument for each endogenous variable and that the data satisfy basic rank conditions. See Nevo and Rosen [34] for details. A different way to see that (5) does not represent a loss of generality is to recognize that one can simply omit all of the endogenous elements of the w-vector without changing the structure of the problem. This would simply exacerbate the existing endogeneity issues with g.

assumption from the conventional multi-market framework for demand estimation, which would require $\rho_{gu} = 0$. A3 nests this condition as a special case. A3 imposes the weaker restriction that the instrument is correlated with the error term in the same direction as the endogenous variable. This is trivially satisfied if the analyst is willing to make an assumption about the direction of correlation between g and u and between z and u. As long as the direction of these correlation coefficients is known, A3 can be satisfied by multiplying the instrument by 1 or -1. Following NR, we refer to instruments that satisfy A3 as "imperfect instrumental variables (IIV)".

To focus attention on the amenity coefficient, δ, it is useful to apply the Frisch-Waugh-Lovell theorem to transform the demand function (5) into a simple bivariate model. First we regress both p_g and g on the vector of covariates, w, and obtain the residuals from these regressions. Let $\tilde{p}_g$ and $\tilde{g}$ denote the two sets of residuals. Using these residuals, we can rewrite the estimator for δ as:

$$\tilde{p}_g = \tilde{g}\delta + u. \tag{6}$$

The probability limits of the OLS and IV estimators for δ are:

$$\delta^{OLS} = \sigma_{\tilde{g}\tilde{p}_g}/\sigma_{\tilde{g}}^2 \quad \text{and} \quad \delta_{z_j}^{IV} = \sigma_{z_j\tilde{p}_g}/\sigma_{z_j\tilde{g}}. \tag{7}$$

If assumptions A1-A3 hold, then δ is bounded by δ^{OLS} and $\delta_{z_j}^{IV}$. NR establish this result as their Lemma 1.[13] We repeat it here for convenience.

Proposition1. Let Assumptions A1-A3 hold.

<u>Case 1</u>: If $\sigma_{z_j\tilde{g}} < 0$, then $min\left\{\delta^{OLS}, \delta_{z_j}^{IV}\right\} \leq \delta \leq max\left\{\delta^{OLS}, \delta_{z_j}^{IV}\right\}$.

[13] NR's lemma 1 is based on a simple bivariate linear model, rather than residuals from multivariate regression. However, this has no real impact on the characterization of the bounds.

Case 2: If $\sigma_{z_j\tilde{g}} > 0$ and $\sigma_{z_ju}, \sigma_{\tilde{g}u} \geq 0$, then $\delta \leq min\left\{\delta^{OLS}, \delta_{z_j}^{IV}\right\}$;

If $\sigma_{z_j\tilde{g}} > 0$ and $\sigma_{z_ju}, \sigma_{\tilde{g}u} \leq 0$, then $\delta \geq max\left\{\delta^{OLS}, \delta_{z_j}^{IV}\right\}$.

Proof: See appendix.

Notice that the direction of correlation between the instrument and the amenity determines whether the IIVs identify one-sided bounds or two-sided bounds. When $\sigma_{z_j\tilde{g}} > 0$ for every IIV, the OLS and IV estimators only bound δ from one side. The other bound is trivially defined by 0 or $-\infty$ under the assumption that demand curves slope down.[14]

We now turn to the problem of measuring consumer surplus. Consider a prospective policy that will produce a non-marginal change in the amenity. Bounds on δ can be used to derive bounds on Marshallian consumer surplus (*MCS*) for the change. Denote the bounds as: $\hat{\delta}_L \leq \delta \leq \hat{\delta}_U$. The corresponding bounds on $\acute{\gamma} = w\gamma$ are defined by $\bar{p}_g - \hat{\delta}_U\bar{g} \leq \acute{\gamma} \leq \bar{p}_g - \hat{\delta}_L\bar{g}$, where $\bar{p}_g$ and $\bar{g}$ represent the baseline levels of $\tilde{p}_g$ and $\tilde{g}$ respectively.[15] Note that $\acute{\gamma}$ is a decreasing function of $\hat{\delta}$ since $\acute{\gamma} = \bar{p}_g - \hat{\delta}\bar{g}$. Therefore, the upper and lower bounds on δ define two demand curves with different slopes and intercepts. Figure 3 provides an example. When δ increases from $\hat{\delta}_L$ to $\hat{\delta}_U$ the demand curve rotates counterclockwise. The two demand curves must share a common intersection at the observed level of the amenity $(\bar{g}, \bar{p}_g)$.

---

[14] NR observe that it is still possible to identify two-sided bounds if there are multiple IIVs and one is both more relevant and more valid than another. Without loss of generality, let z_1 be the more relevant instrument: $\rho_{z_1\tilde{g}} > \rho_{z_2\tilde{g}}$. If we are willing to assume that z_1 is also more valid in the sense that $\rho_{z_1u} < \rho_{z_2u}$, then NR demonstrate it is possible to construct a weighted average of the two instruments that is negatively correlated with the amenity: $h = \phi z_2 - (1-\phi)z_1$, where $\phi = \sigma_{z_1}/(\sigma_{z_1} + \sigma_{z_2})$. If $\sigma_{h\tilde{g}} < 0$, then $\delta_h^{IV} \leq \delta \leq min\{\delta^{OLS}, \delta_{z_1}^{IV}, \delta_{z_2}^{IV}\}$. See NR for a proof and Zhang [40] for additional discussion.

[15] The identification region for each parameter in γ can be obtained using Proposition 3 in NR.

Narrower bounds on δ imply narrower bounds on *MCS*. This can be seen from figure 3. The demand curves in the left panel provide narrower bounds on δ than the demand curves in the right panel. We consider two sets of change in the level of the amenity. First, suppose g decreases from $\bar{g}$ to g_0. Then MCS_1 equals the area -*ghea* and MCS_2 equals -*gheb*, which is smaller in absolute magnitude than MCS_1. Therefore, we have $MCS_1 < MCS < MCS_2$ and the width of these bounds is *abg*. Clearly, the size of area *abg* is smaller in the left panel where we have narrower bounds. Now suppose g increases from $\bar{g}$ to g_1. MCS_1 (represented by *ghfd*) is smaller than MCS_2 (represented by *ghfc*). The bounds for *MCS* are again $MCS_1 < MCS < MCS_2$ and their width is *gdc*. As in case 1, the width is smaller in the left panel than in the right panel. Thus, the bounds on MCS will be tight if the bounds on δ are tight. If the bounds on the slope of the demand curve are one-sided then they only identify an upper bound or a lower bound on the relevant welfare measure.

For the case of two-sided bounds, figure 3 demonstrates that the upper bound on the demand curve does not always provide an upper bound on Marshallian surplus. For a quality decrease the upper bound provides the maximum surplus (|*ghea*|>|*gheb*|) and for a quality increase the upper bound demand function provides the minimum surplus estimate (*ghfd*<*gfhc*). This is important to keep in mind because decision makers evaluating a benefit-cost analysis may want to know the conservative estimate in addition to the bounds on surplus. If the minimum surplus estimate exceeds costs, this may provide decision makers with more confidence than if project feasibility is only indicated by the upper bound estimate of surplus.

4. A Demonstration: Water Quality in Markets for Lakefront Properties

4.1. Data Description

To demonstrate how the partial identification approach works, we use sales data on lakefront properties in Vermont, New Hampshire, and Maine; water quality in the adjacent lakes; and homebuyer demographics. These data were collected from tax records, government agencies, and surveys of homeowners during the mid-1990s (for details see [16,18,27]). While the data are somewhat dated, they have the advantage of having been used in one of the few multi-market hedonic demand studies in the literature [7]. Boyle, Poor, and Taylor's [7] results for Maine provide a baseline to evaluate the economic importance of relaxing the "no stratification" assumption.

The sales data describe residential properties on freshwater lakes and ponds that were sold between 1990 and 1995.[16] Properties include single family residential houses, vacation houses, and unimproved lots. Sale prices were collected from transfer tax records and the property characteristics were obtained from property tax assessment records, both of which are maintained at the individual town offices in New England. The sales data include 230 transactions in Vermont, 518 transactions in New Hampshire, and 851 transactions in Maine. Table 1 provides summary statistics, by state. Each state is subdivided into several regions that real estate agents consider to be distinct markets. There are 20 lakes within 3 regions in Vermont, 53 lakes within 5 regions in New Hampshire, and 37 lakes within 7 regions in Maine. Figure 4 maps the three states, the locations of the lakes, and the market regions.

While our sample sizes may look small relative to recent transaction-based studies with

[16] The period 1990 and 1995 was originally selected because the real estate market for lakefront properties in the study area was stable during this period. Poor et al. [37] tested this stability statistically for the Maine data.

hundreds of thousands of observations, it is important to keep in mind that every data point in our sample describes a transaction for lakefront property. Therefore, we are not attempting to estimate the hedonic price surface for the entire housing market. We are using all sales data from a specialized segment of the market to estimate the portion of the hedonic price surface that applies to lakefront property only. In this sense, our sample sizes allow for statistical precision comparable to other recent studies. For example, the number of lakefront adjacent houses in our sample (1,599) is roughly comparable to the number (2,605) used by Abbott and Klaiber [1] in their recent work on valuing adjacency to urban lakes in desert communities. The somewhat larger sample for Abbott and Klaiber is likely due to their urban application where there would be more residences than would occur for rural lakes.

Lakes in Maine, Vermont, and New Hampshire are generally known for having clear, high-quality water. However, some lakes are threatened by cultural eutrophication that originates from residential development, silviculture, and agricultural activities. The resulting eutrophication reduces water transparency. We use transparency as an indicator of water quality for three reasons. It is highly correlated with other measures of water quality, it is a measure of lake water quality that people observe, and it is the only measure of water quality that is consistently maintained across lakes and over time.[17]

Data on water transparency were provided by the Vermont Department of Environmental Conservation, the Maine Department of Environmental Protection, and the New Hampshire Department of Environmental Services. Transparency is measured using a secchi

[17] Other water chemistry tests are only conducted for specific "hot spots" where eutrophication is known to be particularly problematic [37].

disk that is 8 inches in diameter and alternatively black and white in each quadrant. The disk is lowered into the water and the depth (in meters) that the disk disappears from sight is the measure of transparency. We follow previous studies in using minimum transparency during the summer months as our measure of quality. This is because transparency fluctuates greatly in the spring and fall due to water flows and silt disturbances, and in winter the ice prohibits measurement. Furthermore, since summer months are the time when algal growth is stimulated by long exposure to sunlight, it is the most appropriate time to measure the water's trophic status.

Water transparency varies over the lakes within each market and over time for each individual lake. As eutrophication is a long-term process, the year-to-year variation for each lake is minimal. Most of the variation in water transparency occurs between lakes. Summary statistics are reported in Table 1.

Finally, data on homeowner demographics were obtained from a mail survey of the individuals who purchased each house. Respondents were asked to report the age and employment status of the household head, the total household income, and the number of children in the household.[18] They were also asked questions about their familiarity with the lake. Table 2 provides summary statistics for these variables, by state.

4.2. Research Design

We estimate a hedonic price function for lakefront properties in each market region and then use the results to estimate aggregate demand functions for water quality. It is important to reiterate that our objective is to demonstrate a new approach to the second stage

[18] Households were asked to report their income in narrow bins, similar to the 2000 Census of Population and Housing. We use bin midpoints for the purposes of demand estimation.

of hedonic demand estimation. In contrast, most recent empirical studies have sought to refine the first stage by using quasi-experimental sources of variation in amenities to mitigate omitted variable bias in estimating price function parameters (e.g. [1,15,26,32]). However these studies are ultimately limited by the fact that even well-identified measures of MWTP cannot be used to consistently evaluate non-marginal changes in amenities. In order to develop credible estimates of welfare measures, it is necessary to have credible identification strategies for both the first stage and the second stage. Our application is the converse of most recent quasi-experimental studies. That is, we take a conventional approach to the first stage and focus on improving the credibility of the second stage. Our methods are equally applicable to settings with quasi-experimental identification strategies for first stage price function parameters.

Our interpretation of the first-stage estimates as measures of MWTP is supported by the way our data appear to provide a relatively good approximation to the static equilibrium concept that underlies Rosen's [38] welfare interpretation of the hedonic price function gradient. As noted earlier, all of the housing sales in our data occurred between 1990 and 1995. This was a period of relative stability for housing markets in the Northeast. For example, the Case-Shiller index for Boston changed by less than 4% between January 1990 and December 1995. Likewise, we are unaware of any significant changes in information about water quality or in water quality itself during our study period. The stability of housing markets over our study period approximately satisfies the maintained assumption of stationarity (assumption A1) and allows us to avoid the difficulties with welfare interpretation that can arise when the price function changes over time (e.g. [20,26]).

However, focusing exclusively on lakefront properties suggests a tradeoff between the internal and external validity of our results. The internal validity of our estimate for lakefront property owners' demand for water quality is enhanced by the way our sample avoids the need to control for unobserved variables that systematically differentiate lakefront properties from the rest of the housing market. On the other hand, the resulting demand curve will not necessarily apply to the owners of non-lakefront properties who are excluded from our sample but comprise the majority of the housing market. This tradeoff is similar to the tradeoffs that often arise in quasi-experimental research designs (Parmeter and Pope [36]). For example, hedonic estimates for the MWTP for school quality that leverage boundary discontinuity designs are most applicable to the small subsets of houses that are located near borders between adjacent school zones.

To approximate the sections of the hedonic price surface that apply to each of the 15 markets shown in figure 4, we use simple linear functional forms. This is a departure from theory, which implies that hedonic price functions are generically nonlinear [12]. Further, simulation studies have found that more flexible functional forms such as Box-Cox models tend to provide more accurate estimates for average MWTP when there are no omitted variables (Cropper, Deck, and McConnell [11]) or when it is possible to control for omitted variables using spatial fixed effects (Kuminoff, Parmeter, and Pope [24]). The main reason why we prefer a simpler functional form for the present application is that our sample sizes in the 15 markets are small (from 39 to 254 observations). One of the findings from Kuminoff, Parmeter, and Pope's study is that larger sample sizes are needed to pin down curvature in the price function. We also suspect that linear approximations will be less detrimental when they are applied to relatively small and homogenous segments of the overall housing market. In other words, each linear approximation to a section of the hedonic price function for a group of lakefront properties can be thought of as part of a more flexible "local linear" approximation to the price function for all housing in that market, regardless of proximity to lakes.

Finally, three features of our application make it well suited to partial identification. First, Boyle, Poor, and Taylor [7] find that markets in part of the study area generally have different price functions, suggesting it may be possible to construct instruments from market dummies. Second, the instruments are likely to be imperfect because several of the market areas are in close geographic proximity, suggesting that the cost of migrating between them is likely to be low. While commuting between markets is not extensive in the study area, there are notable exceptions. Table 3 shows county-to-county worker flows for the three states. Virtually no workers commute across state lines. Likewise, within each state, most of working population works in the market area where they live. However, in New Hampshire 22% of workers living in market 1 and 17% of workers living in market 2 commute to work in market 3. One might expect these workers to consider moving closer to their jobs.

More importantly, perhaps, tax records mix sales of "summer houses" with "permanent houses". Purchasers of summer homes may be more likely to search across markets and even all three states, whereas purchasers of permanent homes may be more likely to search in a local area, near their job, or near family and friends. These search patterns could cause the market indicators to be correlated with homeowners' unobserved tastes for water quality.

4.3. Estimation of the Hedonic Price Function

We regress the price of house i in market m on its physical attributes, neighborhood attributes, and environmental amenities:

$$p_{im} = \beta_{0m} + \beta_{1m} WQ_{im} + X_{im}\beta_{2m} + v_{im} \ . \tag{10}$$

The vector X_{im} includes all the physical attributes of the property from table 1. The variable of interest, $WQ_{im} = [LakeSize_{im} \times ln(WT_{im})]$, is an interaction between lake size and the natural log of water transparency (WT). At lower levels of water transparency property owners are expected to pay more for a one meter improvement in transparency. This is consistent with evidence that changes in transparency occurring above four meters are less noticeable than changes below this threshold [40]. In addition, this specification of the quality variable fit the data best in each of the original state-specific studies [16,18,27].

Table 4 reports hedonic estimates for each of the 15 market areas. Most physical characteristics have the expected signs and magnitudes. More lot frontage abutting the lake (*FF*) tends to increase property values. The variable of interest, water quality, significantly increases property values in most markets.

The implicit price for water transparency (*WT*) at each lake can be calculated as follows: $p_{im}^{WT} = \hat{\beta}_{1m}(LakeSize_{im}/WT_{im})$. This becomes the dependent variable used to estimate the willingness-to-pay function. With this in mind, it is important to note the heterogeneity in hedonic gradients across markets; the parameter estimates on the *WQ* variable range from about 2 to 193 in the estimated hedonic functions. This variation is needed to identify bounds on the demand for water transparency. In the second stage, we limit the analysis to markets with enough sales data such that the parameter on WQ is statistically significant. This is done out of concern that including insignificant estimates in the second stage could unnecessarily introduce measurement error into our IIV estimates for demand. Increasing measurement error in the second stage could be problematic given our relatively small sample. In future applications with larger first stage samples it would make sense to include precisely estimated zeros for MWTP in the second stage order to pin down the choke price. We try this as a robustness check (described below) and find that it has little effect on our main results.

4.4. Estimation of the Demand Function

We model the MWTP for water transparency as a function of the transparency level, the square feet of living area and lake frontage, and household demographics including age, income, retirement status, and the number of children in the household. We also control for whether the household had visited the lake previously and whether they have a friend or relative owning a house on the same lake. The estimating equation is specified as:

$$p_i^{WT} = \gamma_0 + \delta WT_i + w_i \gamma_1 + u_i, \tag{11}$$

where w includes all attributes other than water transparency. All else constant, we would expect the households who are in the market for lakefront properties to sort themselves positively over the lakes according to their unobserved tastes for water clarity: $\rho_{WT,u} > 0$. Since lake-based recreation is a major reason for purchasing a lakefront home, it seems reasonable to expect that anyone purchasing a house on a lake would pay careful attention to water clarity and that, all else constant, those with stronger tastes for clarity would choose lakes with less eutrophication.

We develop an instrument, z_1, from a categorical variable indexing the level of water transparency in each market area. With K markets, z_1 takes a value of 1 for the market with lowest average water transparency and a value of K for the market with highest average water transparency.[19] The resulting instrument is similar to the rank-based instruments introduced by Epple and Sieg [14] and used in several applications to modeling household sorting behavior (e.g. [23,39,41]).[20] The intuition is analogous to the conventional multi-market approach to hedonic demand estimation. If households were unable to sort themselves across market areas, then we would expect their unobserved tastes to be independent of the ranking of market areas by water transparency.[21] Following the logic in Bartik [4], we also define a second instrument that interacts the market index with household income: $z_2 = z_1 \times INC$.

If z_1 and z_2 are positively correlated with unobserved tastes for water transparency, violating the standard IV assumption for a multiple market hedonic model, then we can still obtain consistent bounds on the demand curve. Positive correlation seems likely. To the extent that households sort themselves across market areas, we would expect that households with stronger tastes for water clarity would choose areas with clearer lakes ($\rho_{z_1,u} > 0$). Likewise, we would expect the households who can afford to live in areas with clearer lakes will tend to be the wealthier ($\rho_{z_1,u} > 0$). This is consistent with the "income stratification" property that underlies models of sorting equilibria (e.g. see Smith et al. 2004).

[19] In the special case with two markets, our rank-based instrument is equivalent to an indicator for one of the markets.

[20] Epple and Sieg [14] used the income rank of communities as their instruments because they assume that if preferences were homogeneous then communities are stratified by income alone. In our case, we use the water transparency rank of markets as our instruments because it facilitates the analysis of correlation direction between instrument and error term.

[21] In principle, one could use 2SLS with instruments based on indicators for each market area, following the logic of [4,25,35]. However, with multiple indicators, it becomes difficult to develop intuition for the direction of correlation between each indicator and the econometric error term.

Furthermore, wealthier households are more likely to own boats, jet-skis, and other expensive equipment that would allow them to spend more time on the water, increasing their demand for water clarity.[22]

Empirically, both instruments are positively correlated with water transparency. Therefore, under the assumptions of our model, their probability limits will both lie on the same side of the slope of the demand function. We obtain the opposite bound from the trivial assumption that the demand curve is downward sloping.

Panel A of table 5 reports parameter estimates for several specifications of the model. The first three columns report the coefficient on water transparency from a naïve OLS specification, and from conventional IV models using z_1 and z_2 as instruments respectively. Moving from OLS to IV estimation more than doubles the slope coefficient. This is important because the IV models define upper bounds on the slope coefficients. Column 4 reports results from the IIV model using both instruments. Specifically, we report the bounds we obtain from using our assumption of positive taste-based sorting to sign the direction of the correlation between the instruments and error term. Notice that the IIV bounds are informative in the sense that they exclude the OLS point estimate (-718), which lies above the upper bound of the range implied by the IIV model $[-\infty, -3{,}216]$. Column 5 reports 95% confidence intervals on these bounds, and the remaining columns present results from sensitivity analyses that we will discuss below.

Panel B of table 5 reports the results from using the estimated bounds on the water transparency coefficient to calculate bounds on the average consumer surplus from "small"

[22] Previous studies have documented that wealthier people are willing to pay more for improved water quality. For example, Kosenius [21] reported that higher income people are willing to pay more for increased water transparency and decreased occurrences of algal blooms.

and “large” improvements in water clarity. The Maine department of Environmental Protection defines 3 meters of visibility as the threshold below which a lake has significantly compromised water quality. We divide all lakes into groups based on whether they are above or below this threshold. The average visibility for all lakes in our data is 4.7 meters; the average visibility for those with compromised water quality (<3 meters) is 2.1 meters; and the average visibility for lakes that do not have compromised water clarity (>3 meters) is 5.4 meters. We use this information to calculate bounds on the consumer surplus from a relatively small increase in clarity from 4.7 to 5.4 meters and a relatively large increase from 2.1 to 5.4 meters.

Upper bounds on the slopes of the demand curves correspond to upper bounds on the average homeowner’s total gain/loss associated with a given improvement in water transparency. We report our point estimates for these bounds in column (4). Column (5) reports 95% confidence intervals on the point estimates. Comparing the upper bounds on Marshallian consumer surplus in columns (4) and (5) allows us to distinguish between economic uncertainty (column 4) and statistical uncertainty (the difference between columns 4 and 5). For example, our point estimate for the IIV upper bound on the willingness-to-pay for improving visibility from 2.1 meters to 5.4 meters is $26,426. The 95% confidence interval on this bound is somewhat higher ($29,719). Both the economic uncertainty and the statistical uncertainty are important, but most of the width in the bounds is due to economic uncertainty. The statistical component of the uncertainty is only about 10% of the total uncertainty.

Boyle, Poor, and Taylor [7] provide a baseline for comparison. They estimated the demand for water clarity using only the data for Maine and the same functional forms for the price and demand functions.[23] They developed instruments from variables believed to cause price differences across market areas, but not within a single area, such as total population of all towns within each market group divided by the total feet of lake shoreline in the market group region; the total number of lakes in each market; the minimum distance between the center of each market to the nearest business center; and the unemployment rate for each market in the year the property was purchased. Their IV point estimate for the demand function implies a $1,270 consumer surplus for an increase from 4.6 to 5.2 meters. This figure falls within the bounds we obtain for the same improvement in clarity when we repeat estimation of the model in column (4) using data for Maine only (0 to $1,940). Thus, acknowledging the possibility of taste-based sorting implies a more cautious interpretation of the evidence. Yet the bounds are still potentially informative.

4.5. Sensitivity Analysis

We briefly describe three robustness checks. First, we repeat the second stage estimation after including statistically insignificant first stage estimates for MWTP. As noted earlier, our reason for excluding these insignificant estimates from our baseline model is concern that the large confidence intervals may unnecessarily increase measurement error in the second stage. We test three alternative approaches: (i) including statistically insig-

[23] We use a slightly different set of explanatory variables in order to be as consistent as possible across all three states, while using a parsimonious specification. Boyle, Poor, and Taylor [7] included two additional variables in their price function: distance to the nearest city and a dummy indicating whether the property's primary source of water is the lake. In contrast, we added a dummy variable indicting whether the property is bare land. In the demand function, we dropped variables indicating whether the purchaser expected an improvement, decline, or no change in the water clarity at the time the property was purchased. We added demographic variables indicating the purchaser's age, whether he/she is retired, and how many children under 18 in the family. These demographic variables are more widely available and more widely used in second stage hedonic estimation.

nificant point estimates; (ii) including statistically insignificant estimates and replacing the negative and insignificant point estimates with zeros; and (iii) replacing all insignificant point estimates with zeros. Of these three approaches, the last two impose the intuitive restriction that people do not dislike water clarity. The three alternate specifications result in upper bound estimates for WTP that are 1% to 14% lower than our baseline estimates in table 5.[24] It would make sense for future applications with larger first-stage samples to include precisely estimated zeros for MWTP in the second stage. Doing so could help to identify the choke price for the demand curve.

Second, we investigate the sensitivity of our results to the set of demographic demand shifters. The homeowner survey provided a rich set of demographic variables describing households. While it is rarely feasible to conduct an original survey, it is often possible to match a parsimonious set of demographic variables to homebuyers who apply for a mortgage. Publicly available data collected through the Home Mortgage Disclosure Act (HMDA) provide three relevant types of information on mortgage applications: (1) property location including census tract; (2) loan information including amount, date, loan type, property type, loan purpose, and owner-occupancy; and (3) the applicant's ethnicity, race, gender, and gross annual income. Bishop and Timmins [6] develop an algorithm to match HMDA data to individual property transactions.

To mimic a setting in which HMDA data are used to define demand shifters we drop all demographics except for income. Comparing the results to our baseline model measures what is gained in our application by moving from HMDA data to a richer set of demographic characteristics. The results are reported in column (6). Comparing columns (5)

[24] Results are reported in table A1 of the supplemental appendix.

and (6) reveals that significantly expanding the set of demographic characteristics beyond income has a small effect on the results, marginally tightening the upper bound of the 95% confidence interval on the slope of the demand curve from -$2,112 to -$2,202. In contrast, if we repeat estimation of the model in (6) but replace income with age, the upper bound increases substantially to -$617. These results suggest that income is the single most important demographic factor in the demand for water quality. Hence, little information about the demand curve would be lost in this application by relying on HMDA demographics.

Finally, we repeat the estimation using a log-linear specification for the hedonic price function. 95% confidence intervals on the resulting parameter estimates and ranges of values for MCS are reported in column (7) of table 5. Comparing columns (5) and (7) illustrates that moving to a log-linear model increases our estimates for the willingness-to-pay to improve clarity from 4.7 to 5.3 meters by nearly 60%. While the log-linear model gives much larger estimates of the WTP for improved water clarity, this upper bound can still be informative for policy as it provides an extreme outer limit of benefits or losses. If a policy does not pass a benefit-cost test with the log-linear specification it is fairly strong evidence against a project.

The large difference between the linear and log-linear models appears to be driven by the right tail of the housing price distribution. If we repeat the estimation but exclude the top 5% of housing values, then the log-linear willingness to pay to improve clarity from 4.7 to 5.3 meters is $2,929 compared to $2,629 for the corresponding linear model. While we do not actually want to exclude high price houses due to concerns about sample selection, this example illustrates that the log-linear estimates are more sensitive to outliers. This is expected given the small sizes of our first-stage samples (from 39 to 254 observations per market). Kuminoff, Parmeter, and Pope [24] find that larger sample sizes are needed to pin down curvature in hedonic price functions. With sample sizes of 2000, they find little difference between the performance of linear, log-linear, and log-log specifications, and improvement from moving to more flexible quadratic and Box-Cox specifications. Hence it would make sense for future applications based on larger first-stage samples to use more flexible functional forms.

Even with much larger samples, we would still expect there to be some error in estimating MWTP due to functional form misspecification. This follows from the fact that most utility functions do not generate closed form expressions for the equilibrium price function. If we think of functional form uncertainty as being a component of the economic uncertainty associated with demand function estimation, then our bounds on welfare measures would be defined by the union of the two sets of bounds. That is, not considering the model uncertainty as a component of economic uncertainty could cause an analyst to understate potential gains or losses from changes in quality in this application. While we would expect model uncertainty to be less important in future applications that have larger first-stage samples, it could also be useful to investigate this by extending the simulation framework from Cropper, Deck, McConnell [11] and Kuminoff, Parmeter, and Pope [24] to calculate the probability of Type I and type II errors in using IIV estimates to bound the willingness to pay for large changes in amenities.

5. Conclusions

Taste-based sorting is widely believed to confound the multi-market hedonic approach to reduced-form estimation of demand curves for environmental amenities. This perception stems from the literature's historical focus on point identification. We have shown that it is possible to identify bounds on demand curves under mild assumptions about the nature of sorting behavior. Bounds on demand curves translate into ranges of partial equilibrium welfare measures for non-marginal changes in amenities. Even one-sided bounds can be informative for benefit-cost analysis if, for example, the costs of a regulation exceed an estimated upper bound on benefits. In this sense, we were able to provide an example of informative bounds on the willingness to pay using an application with readily available data; reducing eutrophication of lakes in Maine, New Hampshire, and Vermont.

Ideally, future applications would combine the partial identification approach with richer micro data on housing transactions and best practices for mitigating omitted variable bias in the estimation of first stage hedonic price functions. While demand estimation can be performed using as few as two markets, more markets can provide more information. It is important to keep in mind that the relevant definition of a market for hedonic demand estimation is a spatiotemporal interval that represents equilibrium such that the price function is stable. Recent research has shown that large shocks to amenities, wealth, and information can cause hedonic price functions for housing to change over time [20,24,26]. In principle, this would enable one to identify bounds on amenity demand curves using data on different hedonic price functions describing the same spatial area at different points in time, relaxing the exogeneity assumption in Beron, Murdoch, and Thayer [5] and similar studies.

We also demonstrated that the nature of taste-based sorting determines whether the standard "multiple-market" instrumental variables strategy identifies one-sided or two-sided bounds on the demand for an amenity. If consumers with stronger tastes for the amenity tend to locate in higher amenity areas, then the standard instruments identify an upper bound *or* a lower bound. In contrast, the standard instruments identify an upper bound *and* a lower bound if consumers with stronger tastes for the amenity tend to locate in lower amenity areas. At least three mechanisms have the potential to generate this counterintuitive stratification pattern. First, it could arise if the amenity is negatively correlated with another attribute that plays a larger role in consumer decisions (e.g. preferences for air quality are positive correlated with job skill and higher-skill workers tend to move to cities where the return to skill and pollution are both higher). Second, it could arise from imperfect information and learning (e.g. households who experience an unexpected contamination of private well water subsequently learn about the risks involved with private wells and develop stronger tastes for connection to public water systems). Finally, it could arise from social interactions and feedback effects (e.g. households who prefer quiet move to the suburbs, which induces planners to build noisy freeways connecting suburbs to the city).

Our focus on the housing market was meant to be demonstrative. The partial identification hedonic framework we presented is generalizable to a wide variety of settings where the purchase of a differentiated private good conveys a quasi-public attribute. The estimation process has the benefit of simplicity. It only requires estimating the conventional reduced-form OLS and IV specifications for demand and then using basic knowledge of sorting behavior to interpret the point estimates on demand parameters as bounds. Thus, Nevo and Rosen's [34] recent advances in the econometrics of partial identification of linear IV models provides an opportunity to rehabilitate Rosen's [38] original proposal for a simple reduced-form approach to hedonic demand estimation, and this can be used to improve future research on the demand for nonmarket goods and services.

Appendix: Proof of proposition 1

Proof. Suppose $\rho_{gu} \geq 0$ and $\rho_{z_j u} \geq 0$. Then we have:

$$\rho_{gu} \geq 0 \quad \Leftrightarrow \quad \sigma_{gu} \geq 0 \quad \Leftrightarrow \quad \sigma_{g(\tilde{p}_g - \delta\tilde{g})} \geq 0$$

$$\Leftrightarrow \quad \sigma_{g\tilde{p}_g} - \delta\sigma_{g\tilde{g}} \geq 0 \quad \Leftrightarrow \quad \delta \leq \sigma_{g\tilde{p}_g}/\sigma_{g\tilde{g}} \equiv \delta^{OLS}.$$

The third term uses $\tilde{p}_g = \delta\tilde{g} + u$, where $\tilde{p}_g$ and $\tilde{g}$ denote the residuals from regressing p_g and g on the vector of covariates. In the last term, $\sigma_{g\tilde{p}_g}/\sigma_{g\tilde{g}} = \delta^{OLS}$ because

$$\sigma_{gu} = 0 \quad \Leftrightarrow \quad \sigma_{g(\tilde{p}_g - \delta^{OLS}\tilde{g})} = 0 \quad \Leftrightarrow \quad \delta^{OLS} = \sigma_{g\tilde{p}_g}/\sigma_{g\tilde{g}}.$$

Similarly,

$$\rho_{z_j u} \geq 0 \quad \Leftrightarrow \quad \sigma_{z_j u} \geq 0 \quad \Leftrightarrow \quad \sigma_{z_j(\tilde{p}_g - \delta\tilde{g})} \geq 0 \quad \Leftrightarrow \quad \delta\sigma_{z_j\tilde{g}} \leq \sigma_{z_j\tilde{p}_g}.$$

If $\sigma_{z_j\tilde{g}} > 0$, then $\delta \leq \sigma_{z_j\tilde{p}_g}/\sigma_{z_j\tilde{g}} \equiv \delta_{z_j}^{IV}$. Alternatively, if $\sigma_{z_j\tilde{g}} < 0$, then $\delta \geq \sigma_{z_j\tilde{p}_g}/\sigma_{z_j\tilde{g}} \equiv \delta_{z_j}^{IV}$. The proposition is completed by combining these inequalities with symmetric reasoning for the case where $\rho_{gu} \leq 0$ and $\rho_{z_j u} \leq 0$.

References

[1] Abbott, J.K. and H.A. Klaiber. 2013. "The Value of Water as an Urban Club Good: A Matching Approach to Community-Provided Lakes." Forthcoming in *Journal of Environmental Economics and Management.*

[2] Bajari, P., and C. L. Benkard. 2005. "Demand estimation with heterogeneous consumers and unobserved product characteristics: a hedonic approach", *Journal of Political Economy*, 113(6), 1239-1276.

[3] Bajari, P., and M. E. Kahn. 2005. "Estimating housing demand with an application to explaining racial segregation in cities", *Journal of Business and Economic Statistics*, 23(1), 20-33.

[4] Bartik, T. J. 1987. "The estimation of demand parameters in hedonic price models", *Journal of Political Economy*, 95(1), 81-88.

[5] Beron, K., J. Murdoch, and M. Thayer. 2001. "The Benefits of Visibility Improvement: New Evidence from the Los Angeles Metropolitan Area." *Journal of Real Estate Finance and Economics*. 22(2/3): 319-337.

[6] Bishop, K., and C. Timmins. 2008. "Simple, consistent estimation of the marginal willingness to pay function: recovering Rosen's second stage without instrumental variables", *working paper*.

[7] Boyle, K. J., P. J. Poor, and L. O. Taylor. 1999. "Estimating the demand for protecting freshwater lakes from eutrophication", *American Journal of Agricultural Economics*, 81(5), 1118-1122.

[8] Brown, J.N. and H.S. Rosen. 1982. "On the Estimation of Structural Hedonic Price Models."*Econometrica*, 50(3): 765-68.

[9] Chattopadhyay, S. 1999. "Estimating the demand for air quality: new evidence based on the Chicago housing market", *Land Economics*, 75(1), 22-38.

[10] Chen, Y. and S. S. Rosenthal. 2008. "Local Amenities and Life-cycle Migration: Do People Move for Jobs or Fun?" *Journal of Urban Economics*. 64(4): 519-537.

[11] Cropper, M.L., L.B. Deck, and K.E. McConnell. 1988. "On the Choice of Functional Form for Hedonic Price Functions," *Review of Economics and Statistics* 70, 668-675.

[12] Ekeland, I., J.J. Heckman, and L. Nesheim. 2004. "Identification and Estimation of Hedonic Models," *Journal of Political Economy* 112, S60-S109.

[13] Epple, D. 1987. "Hedonic prices and implicit markets: estimating demand and supply functions for differentiated products", *Journal of Political Economy*, 95(1), 59-80.

[14] Epple, D., and H. Sieg. 1999. "Estimating equilibrium models of local jurisdictions", *Journal of Political Economy*, 107(4), 645-681.

[15] Gamper-Rabindran, S., and C. Timmins. 2013. "Does Cleanup of Hazardous Waste Sites Raise Housing Values? Evidence of Spatially Localized Benefits." Forthcoming in *Journal of Environmental Economics and Management.*

[16] Gibbs, J. P., Halstead, J. M., Boyle, K. J., and Huang J. C. 2002. "An hedonic analysis of the effects of lake water clarity on New Hampshire lakefront properties", *Agricultural and Resource Economics Review*, 31(1), 39-46.

[17] Graves P.E. and D.M. Waldman. 1991. "Multimarket Amenity Compensation and the Behavior of the Elderly." *American Economic Review* 81(5): 1374-1381.

[18] Hsu, T. I. 2000. "Hedonic study of the effects of lake-water clarity and aquatic plants on lakefront property prices in Vermont", *Master's Thesis, University of Maine*.

[19] Klaiber, H. A., and D. J. Phaneuf. 2010. "Valuing open space in a residential sorting model of the Twin Cities". *Journal of Environmental Economics and Management*, 60(2), 57-77.

[20] Klaiber, H. A., and V. K. Smith. 2013. "Quasi Experiments, Hedonic Models, and Estimating Trade-offs for Local Amenities". *Land Economics*. 89(3): 413-431.

[21] Kosenius, A. K. 2010. "Heterogeneous preferences for water quality attributes: The Case of eutrophication in the Gulf of Finland, the Baltic Sea". *Journal of Ecological Economics*, 69(3), 528-538.

[22] Krupka, D.J. 2009. "Location-Specific Human Capital, Location Choice, and Amenity Demand." *Journal of Regional Science*. 49(5): 833-854.

[23] Kuminoff, N.V. 2009. "Decomposing the Structural Identification of Nonmarket Values." *Journal of Environmental Economics and Management*, 57(2): 123-139.

[24] Kuminoff, N.V., C. F. Parmeter and J.C. Pope. 2010. "Which Hedonic Models Can We Trust to Recover the Marginal Willingness to Pay for Environmental Amenities". *Journal of Environmental Economics and Management* 60(3): 145-160.

[25] Kuminoff, N.V., and J.C. Pope. 2012. "A Novel Approach to Identifying Hedonic Demand Parameters". *Economics Letters* 116(3): 374-376.

[26] Kuminoff, N.V. and J.C. Pope. 2014. "Do "Capitalization Effects" for Public Goods Reveal the Public's Willingness to Pay?" *International Economic Review* 55(4): 1227-1250.

[27] Lawson, S. R. 1997. "Estimating the benefits of water quality in Maine's lakes: a hedonic property value model", *Master's Thesis, University of Maine*.

[28] Leamer, E.E. 1983. "Let's Take the Con Out of Econometrics." *American Economic Review*, 73(1): 31-43.

[29] Linneman, P. 1980. "Some empirical results on the nature of the hedonic price function for the urban housing market", *Journal of Urban Economics*, 8(1), 47-68.

[30] Manski C. 2007. *Identification for Prediction and Decision.* Harvard University Press. Cambridge, MA.

[31] McConnell, K.E. and T.T. Phipps. 1987. "Identification of Preference Parameters in Hedonic Models: Consumer Demands with Nonlinear Budgets." *Journal of Urban Economics*. 22: 35-52.

[32] Muehlenbachs, L., E. Spiller, and C. Timmins. 2012. "Shale Gas Development and Property Values: Differences across Drinking Water Sources." *NBER Working Paper* #18390.

[33] Mueser, P.R. and P.E. Graves. 1995. "Examining the Role of Economic Opportunity and Amenities in Explaining Population Redistribution." *Journal of Urban Economics* 37: 176-200.

[34] Nevo, A., and A. M. Rosen. 2012. "Identification with imperfect instruments", *The Review of Economics and Statistics*, 94(3): 659-671.

[35] Palmquist, R. B. 1984. “Estimating the demand for the characteristics of housing”, *The Review of Economics and Statistics*, 66(3), 394-404.

[36] Parmeter, C. F. and J. C. Pope. 2013. “Quasi-Experiments and Hedonic Property Value Methods,” in J. A. List and M. K. Price, eds. *Handbook on Experimental Economics and the Environment* (Northampton, MA: Edward Elgar Publishing), 3-66.

[37] Poor, P. J., K. J. Boyle, L. O. Taylor, and R. Bouchard. 2001. “Objective versus subjective measures of water clarity in hedonic property value models”, *Land Economics*, 77(4), 482-493.

[38] Rosen, S. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy* 82, 34-55.

[39] Sieg, H., V. K. Smith, H. S. Banzhaf, and R. P. Walsh. 2004. “Estimating the general equilibrium benefits of large changes in spatially delineated public goods”, *International Economic Review*, 45(4), 1047-1077.

[40] Smeltzer, E., and S. Heiskary. 1990. “Analysis and application of lake survey data”, *Lake and Reservoir Management*, 6(1), 109-118.

[41] Smith, V. K., H. Sieg, H. S. Banzhaf, and R. P. 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), 559-584

[42] Walsh, R. 2007. “Endogenous Open Space Amenities in a Locational Equilibrium.” *Journal of Urban Economics*, 61(2), 319-344.

[43] Witte, A. D., H. J. Sumka, and H. Erekson. 1979. “An estimate of a structural hedonic price model of the housing market: an application of Rosen's theory of implicit markets”, *Econometrica*, 47(5), 1151-1173.

[44] Zabel, J. E., and K. A. Kiel. 2000. “Estimating the demand for air quality in four U.S. cities”, *Land Economics*, 76(2), 174-194.

[45] Zhang, C. 2012. “Essays on the Use of Hedonic Price Models to Measure Welfare for Quality Changes in Public Goods.” *PhD. Dissertation*, Virginia Tech.

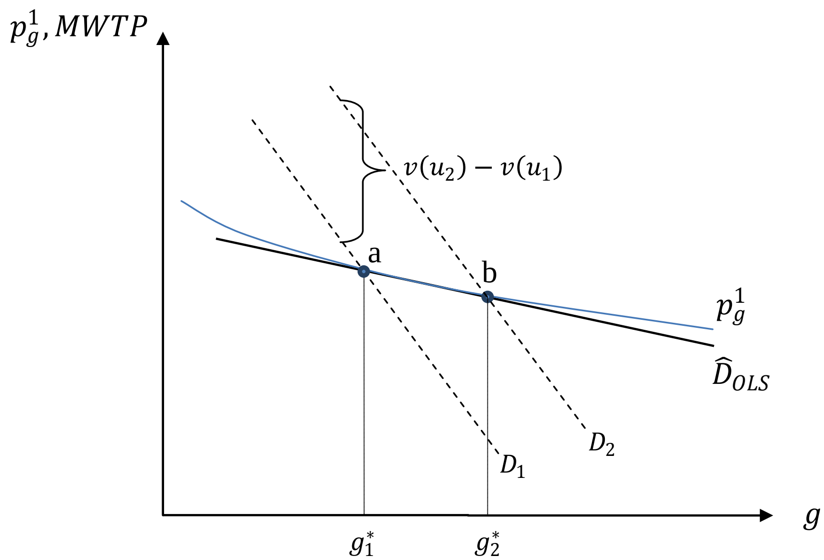

A. Endogeneity due to taste-based sorting within a market

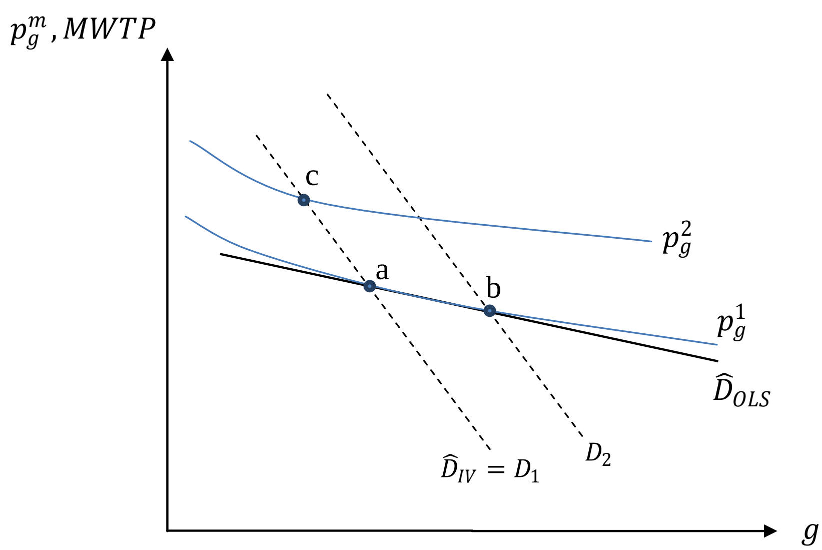

B. Point Identification from random assignment of "D_1-type" households across markets

Figure 1. An Identification Problem with Hedonic Demand Estimation

Note: In market #1, two households' inverse demand curves for an amenity (D_1 and D_2) intersect the marginal price function for that amenity (p_g^1). The analyst observes the utility-maximizing quantities and implicit prices selected by each household (**a** and **b**). Panel A illustrates that a line connecting the two points understates the slopes of individual demand curves because the household with stronger unobserved tastes ($v_2 > v_1$) selects a larger quantity of g. Panel B illustrates that D_1 is identified if an identical household facing a different price function in metro area #2 is observed at a different point on their shared demand curve, such as point **c**.

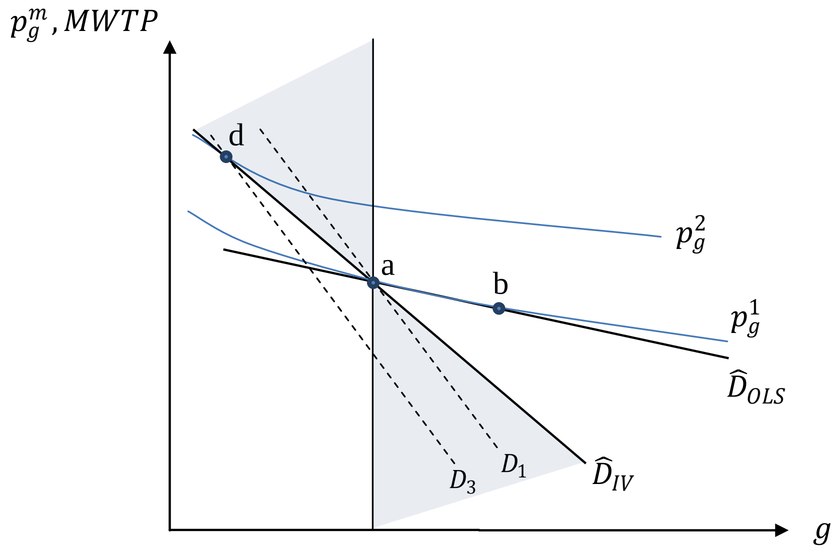

A. Bounds on D_1 under positive taste-based sorting across markets

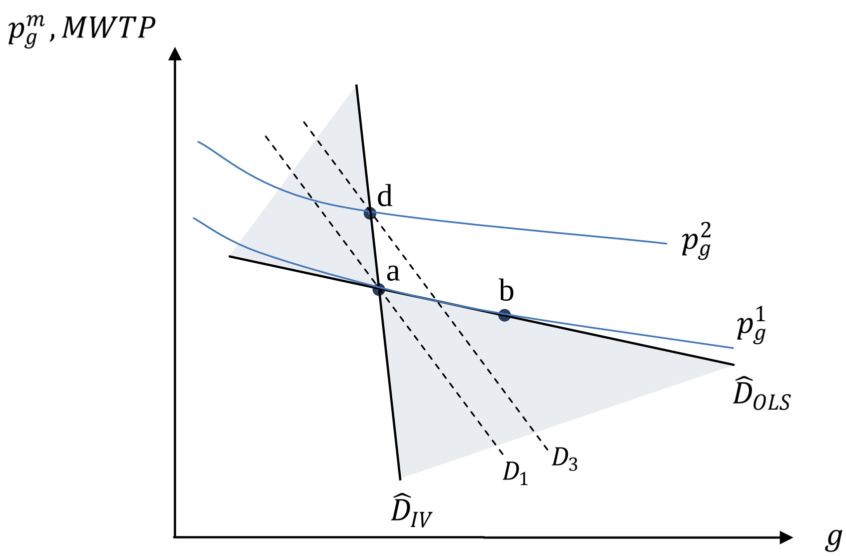

B. Bounds on D_1 under negative taste-based sorting across markets

Figure 2. Using Multiple Markets to Bound Demand Curves

Note: In market #1, a household with inverse demand curve D_1 is observed at point **a**. D_1 is not point identified if households are systematically sorted across markets according to unobserved features of their tastes. Fortunately, D_1 may still be partially identified. Consider an observationally equivalent household in market #2 with demand curve D_3 who maximizes utility by consuming at point **d**. In panel A, $D_1 > D_3$ and D_1 is bounded from above by the IV estimator (connecting points **d** and **a**). It is bounded from below by a perfectly inelastic demand curve through **a**. Panel B illustrates the opposite case where $D_1 < D_3$. In this case, D_1 is bounded by the area #1 OLS estimator and the dual-area IV estimator.

Figure 3. Bounds on the Willingness to Pay for Changes in Environmental Quality

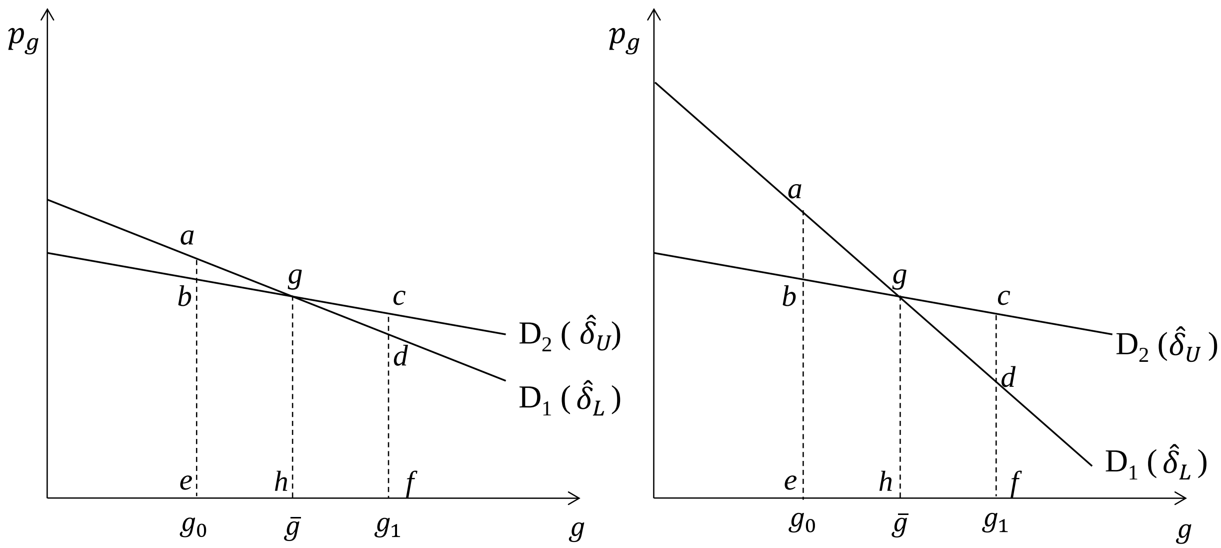

Change	MCS_1	MCS_2	MCS_2 - MCS_1
$\bar{g}$ to g_0	*-ghea*	*-gheb*	*abg*
$\bar{g}$ to g_1	*ghfd*	*ghfc*	*gdc*

Note: The figure illustrates demand curves for an amenity (D_1 and D_2) based on upper and lower bounds on the slope coefficient in the inverse demand function. The bounds are wider in the figure on the right. The table shows Marshallian welfare measures for quality changes relative to a baseline of $\bar{g}$.

Figure 4. Lakes and Market Areas in Three States

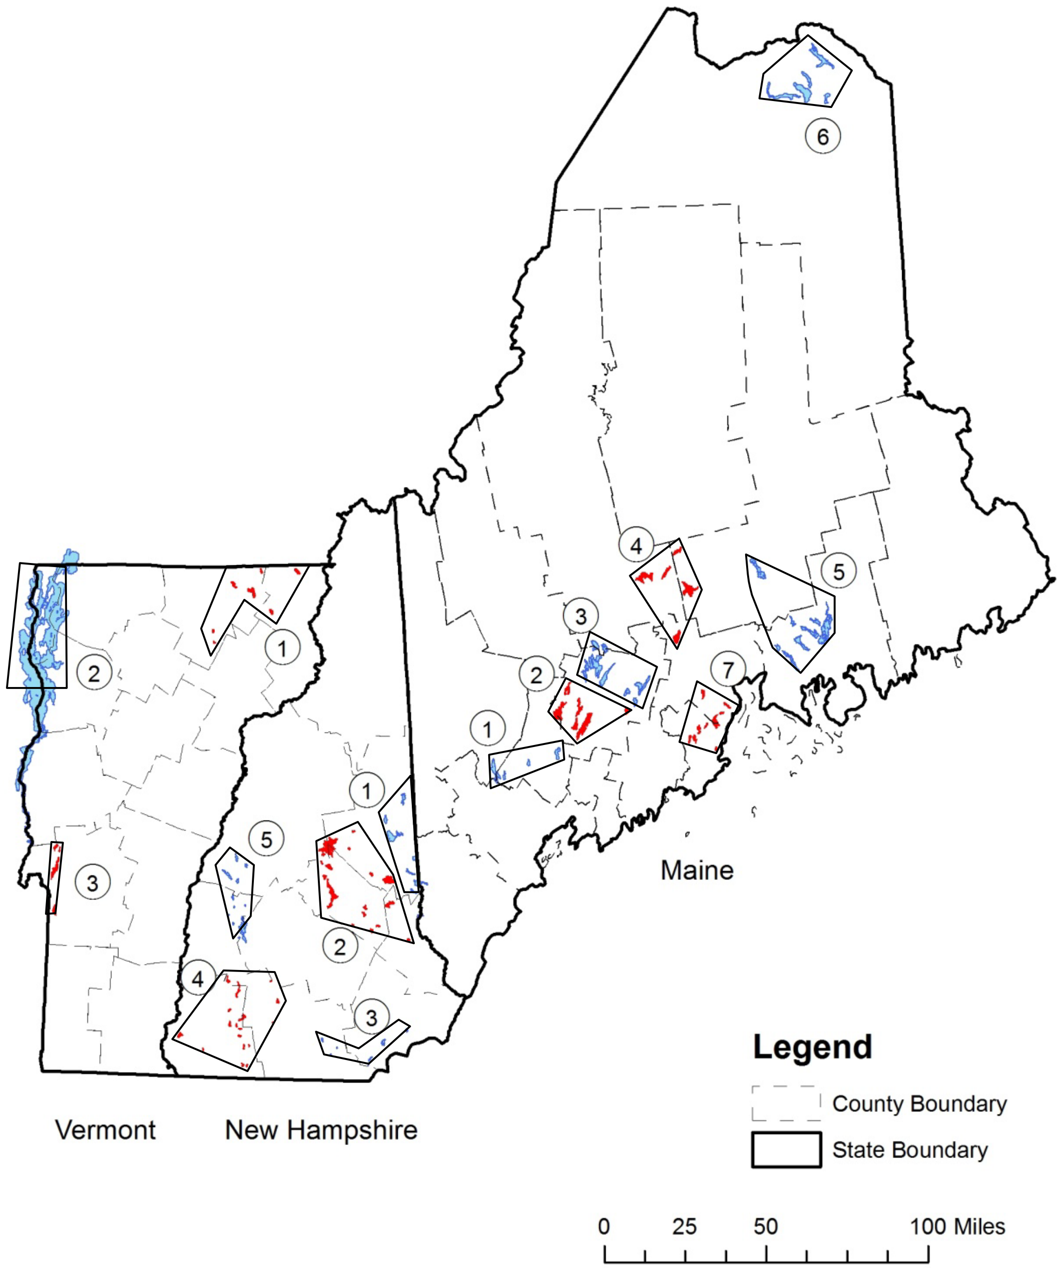

Note: The numbered polygons correspond to distinct market areas that real estate agents treat as distinct markets. The shaded areas within each polygon are lakes. Red and blue shading is used to help distinguish lakes in market areas that are in close proximity.

Table 1. Summary Statistics for Property Transactions

Variable	Definition	Mean	Std. Dev.	Min.	Max.
	Panel A: Vermont (N=230)				
P	sale price of the property ($1995)	99,034	61,314	4,000	340,000
BARE	0,1 = unimproved land	0.18	0.39	0	1
SQFT	total living area (square feet)	810	578	0	3,560
LOT	lot size (acres)	1.01	1.84	0.07	11.91
HEAT	0,1 = central heating system	0.53	0.50	0	1
FULLBATH	0,1 = a full bathroom	0.68	0.47	0	1
FF	total lot frontage abutting the water (feet)	133	117	15	1,201
LAKESIZE	surface area of lake (acres)	1,575	841	54	2,795
WT	water transparency (meters)	5.26	1.87	1.9	9.5
	Panel B: Maine (N=851)				
P	sale price of the property ($1995)	71,536	57,155	400	500,000
BARE	0,1 = unimproved land	0.25	0.43	0	1
SQFT	total living area (square feet)	715	614	0	4,128
LOT	lot size (acres)	1.37	2.57	0.04	20
HEAT	0,1 = central heating system	0.46	0.50	0	1
FULLBATH	0,1 = a full bathroom	0.59	0.49	0	1
FF	total lot frontage abutting the water (feet)	154	140	10	1,800
LAKESIZE	surface area of lake (acres)	3,515	2,428	171	8,239
WT	water transparency (meters)	4.15	1.98	0.3	9.4
	Panel C: New Hampshire (N=518)				
P	sale price of the property ($1995)	159,299	109,833	12,500	815,254
SQFT	total living area (square feet)	1,127	724	107	6,532
FF	total lot frontage abutting the water (feet)	136	188	5	3,395
LAKESIZE	surface area of lake (acres)	1,241	1,508	31	9,091
WT	water transparency (meters)	4.79	2.05	0.7	11.5
DENSITY	density (lots/1000 ft. of lake frontage)	8.51	3.10	1	20
TAX	property tax rate in year of purchase	22.53	7.80	8.2	41
AGE	age of house (years)	39.47	25.84	0	181
PLUMB	type of plumbing	3.89	0.50	0	4
DIST	distance to the nearest large town	12.36	7.90	0	33

Note: These data were collected from tax records, government agencies, and surveys of homeowners during the mid-1990s. The available variables for New Hampshire do not overlap entirely with Vermont and Maine. For details of the original data collection process see Lawson [23], Hsu [16], and Gibbs et al. [14].

Table 2. Summary Statistics for Demographic Characteristics of Homebuyers

Variable	Description	Mean	Std. Dev.	Min.	Max.
	Panel A: Vermont (N=95)				
RESAGE	age of the mail survey respondent	47	11	25	72
INC	total after-tax household income	89,579	69,848	17,500	350,000
RETIRED	0,1 = respondent is fully retired	0.29	0.46	0	1
KIDS	total number of people under 18-year old in the household	3.01	1.45	1	9
VISIT	0,1 = visited the lake before purchasing the property	0.97	0.18	0	1
FRIEND	friends or relatives own property on the lake	0.45	0.50	0	1
	Panel B: Maine (N=240)				
RESAGE	age of the mail survey respondent	45	11	22	74
INC	total after-tax household income	77,979	50,810	7,500	212,500
RETIRED	0,1 = respondent is fully retired	0.10	0.29	0	1
KIDS	total number of people under 18-year old in the household	0.77	1.04	0	4
VISIT	0,1 = visited the lake before purchasing the property	0.14	0.35	0	1
FRIEND	friends or relatives own property on the lake	0.50	0.50	0	1
	Panel C: New Hampshire (N = 203)				
RESAGE	age of the mail survey respondent	46	11	22	79
INC	total after-tax household income	130,283	95,849	22,500	350,000
RETIRED	0,1 = respondent is fully retired	0.27	0.45	0	1
KIDS	total number of people under 18-year old in the household	0.85	1.07	0	4
VISIT	0,1 = visited the lake before purchasing the property	0.89	0.32	0	1
FRIEND	friends or relatives own property on the lake	0.54	0.50	0	1

Table 3. County-to-County Worker Flows in Maine, Vermont, and New Hampshire

Home Market	Work Market														
	VT 1	VT 2	VT 3	HN 1	NH 2	NH 3	NH4	NH5	ME1	ME2	ME3	ME4	ME5	ME6	ME7
VT 1 (Orleans/Essex)	**77.3**	1.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
VT 2 (Franklin/Grand Isle/Chittenden)	0.1	**94.9**	0.1	0.0	0.0	0.1	0.0	2.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0
VT 3 (Rutland)	0.0	0.7	**86.4**	0.0	0.0	0.0	0.1	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0
NH 1 (Carroll)	0.0	0.0	0.0	**65.3**	*	21.8	1.6	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0
HN 2 (Carroll/Belknap/Strafford)	0.0	0.0	0.0	*	**68.1**	17.4	2.0	1.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
HN 3 (Hillsborough/Rockingham)	0.0	0.0	0.0	1.3	1.4	**74.7**	*	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
HN 4 (Cheshire/Hillsborough)	0.0	0.0	0.0	0.2	0.3	*	**74.0**	0.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0
NH 5 (Sullivan/Grafton)	0.0	0.0	0.0	0.4	3.1	1.2	1.0	**81.9**	0.0	0.0	0.0	0.0	0.0	0.0	0.0
ME 1 (Androscoggin)	0.0	0.0	0.0	0.0	0.0	0.2	0.1	0.0	**74.0**	2.7	2.7	0.2	0.1	0.0	0.0
ME 2 (Kennebec)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	4.1	**80.4**	*	6.1	0.8	0.0	1.5
ME 3 (Kennebec)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	4.1	*	**80.4**	6.1	0.8	0.0	1.5
ME 4 (Penobscot/Somerset/Waldo)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	5.3	5.3	**86.8**	*	0.3	*
ME 5 (Hancock/Penobscot)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.6	0.6	*	**93.9**	0.4	1.1
ME 6 (Aroostook)	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.1	0.0	0.0	2.0	2.0	**96.1**	0.0
ME 7 (Waldo/Knox)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	4.6	4.6	*	8.5	0.0	**80.5**

Note: The table uses Census data on county-by-county worker flow files for 2000 to approximate commuting patterns between defined markets. Markets numbers correspond to the regions shown in figure 4. Counties are shown in parenthesis next to each market. In cases where a market overlaps multiple counties, we report commuting patterns for the counties aggregated. For example, VT1 overlaps parts of Orleans and Essex counties. Each number in the table defines the percentage of a home market's working population that works in the corresponding work market. For example, 77.3% of Orleans/Essex's working population works in the Orleans/Essex County. Asterisks indicate that the statistic was not computed because the home and work markets share a subset of the same counties.

Table 4. Estimates for Hedonic Price Function Parameters

	VT1	VT2	VT3	ME1	ME2	ME3	ME4	ME5	ME6	ME7	NH1	NH2	NH3	NH4	NH5
WQ	**11.4*** **	**10.6*** **	**2.0**	**7.5*** **	**-0.9**	**1.6*** **	**11.4*** **	**-0.1**	**-0.4**	**35.0*** **	**4.6** **	**5.9*** **	**193.3** **	**124.7*** **	**23.8*** **
	(2.1)	(4.0)	(3.7)	(1.4)	(2.6)	(0.4)	(3.4)	(1.2)	(0.5)	(13.5)	(2.0)	(2.0)	(83.4)	(21.1)	(6.6)
SQFT	33.9***	60.2***	71.4***	22.0**	31.7***	48.1***	51.3***	48.5***	35.1***	59.7***	41.1***	50.5***	43.0***	41.0***	109.6***
	(9.1)	(10.7)	(16.5)	(10.5)	(8.1)	(5.0)	(13.4)	(6.0)	(4.3)	(19.3)	(11.6)	(6.9)	(9.0)	(13.5)	(20.5)
FF	127.5***	-4.6	41.1	256.3***	115.7**	57.7**	-46.4***	83.7***	65.3***	126.8***	21.1	107.0*	165.2**	130.3	211.7*
	(39.8)	(45.1)	(95.2)	(67.3)	(47.2)	(23.5)	(17.9)	(19.6)	(17.6)	(35.4)	(15.1)	(57.6)	(82.7)	(88.1)	(119.6)
BARE	-11,678.8	36,537.5	36,442.1	-4,167.8	17,346.2	11,208.1	-10,613.8	11,156.4	11,038.3***	27,806.3					
	(9,894.3)	(22,544.0)	(32,707.5)	(22,254.6)	(14,162.6)	(8,692.5)	(11,828.5)	(7,369.5)	(4,213.6)	(30,125.4)					
LOT	-264.7	-214.8	10,448.0***	-3,830.0	-2,607.2	311.0	15,509.8***	416.1	-1,268.4	-6,322.7***					
	(1,599.1)	(3,488.9)	(3,995.3)	(3,529.4)	(1,830.4)	(1,008.9)	(5,576.3)	(902.6)	(1,360.2)	(2,288.6)					
HEAT	14,839.6*	45,282.2***	6,761.5	30,800.8**	17,674.9**	23,995.8***	2,972.4	19,169.9***	1,348.6	-25,475.5					
	(7,673.5)	(13,366.2)	(19,654.3)	(12,399.2)	(8,976.8)	(5,420.3)	(8,525.7)	(5,644.3)	(4,020.2)	(31,573.7)					
FULLBATH	4,606.4	-4,959.3	24,483.8	19,006.4	24,360.0**	15,523.0**	2,233.8	15,556.3**	15,141.1***	59,507.4					
	(7,172.6)	(20,106.3)	(27,951.9)	(20,409.5)	(11,136.9)	(7,511.8)	(8,787.0)	(6,315.2)	(3,562.2)	(41,584.2)					
AGE											-105.9	-453.0**	-1,746.0**	-1,817.1**	2,115.3**
											(687.2)	(221.4)	(847.1)	(858.0)	(1,075.1)
AGE2											-0.4	-0.0	10.6	12.2*	-15.3**
											(8.3)	(0.0)	(8.7)	(7.1)	(7.8)
PLUMB											23,074.6	6,112.2	2,025.8	-2,371.5	9,646.8
											(17,504.2)	(8,222.3)	(14,863.5)	(13,126.9)	(18,497.1)
DIST											-2,422.7***	-601.3	-714.3	201.8	-216.0
											(740.6)	(877.3)	(1,330.4)	(1,480.8)	(3,525.8)
DENSTY											-1,855.2	-3,648.2*	943.6	-7,663.4**	-5,648.3
											(1,599.4)	(2,005.2)	(2,113.9)	(3,078.2)	(4,220.4)
TAXRT											-3,214.0***	-1,418.7*	304.1	368.1	-4,634.2*
											(816.6)	(726.2)	(1,077.7)	(1,227.4)	(2,451.4)
Intercept	10,294.3	3,869.0	-3,683.3	-19,755.4	10,179.4	2,763.3	2,431.3	2,216.8	-4,602.5	-14,566.2	110,645.4	125,738.5**	59,451.5	136,745.5**	54,573.0
	(8,980.9)	(23,612.9)	(30,590.2)	(21,691.0)	(13,738.2)	(7,746.1)	(8,993.7)	(7,757.4)	(4,812.5)	(26,190.7)	(76,556.8)	(49,175.7)	(67,418.5)	(68,729.7)	(120,199.9)
N	60	99	71	83	105	254	39	164	150	56	110	170	68	73	97
Adjusted R^2	0.63	0.51	0.42	0.54	0.44	0.59	0.80	0.66	0.67	0.62	0.43	0.50	0.62	0.67	0.68

Note: *** p<0.01, ** p<0.05, * p<0.1

Table 5: Estimates for Demand Parameters and Welfare Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
model	OLS	IV1	IV2	IIV	IIV 95% CI	IIV 95% CI	IIV 95% CI
dependent variable in price function	price	price	price	price	price	price	ln(price)
demographics in demand function	all	all	all	all	all	income only	all
Panel A: Estimates for slope (δ) of the demand curve (standard errors) [bounds]							
	-718***	-2154***	-3216***	[-∞, -3,216]	[-∞, -2,202]	[-∞, -2,112]	[-∞, -6,045]
	(144)	(362)	(516)				
Panel B: Upper bound on the range of consistent IIV estimates for consumer surplus from water quality improvements							
clarity improves from 4.7 to 5.4 meters				$2,679	$2,903	$2,944	$4,593
clarity improves from 2.1 to 5.4 meters				$26,426	$29,719	$29,157	$67,154

Note: In panel A, columns (1)-(3) report point estimates and standard errors for the slopes of the demand curves from OLS and IV estimation. Column (4) reports the range of estimates from IIV estimation. Columns (5)-(7) report the ranges of estimates implied by 95% confidence intervals on the endpoints of the ranges from IIV estimation. Panel B reports upper bounds on the ranges of consistent IIV estimates for consumer surplus. The lower bounds are zero. Columns (1) through (5) and (7) include the following demographic variables in the demand function: income, age, retirement status, number of children in the household, and indicators for whether the household had visited the lake previously and whether they have a friend or relative owning a house on the same lake. See the text for additional details.

Table A1: Sensitivity of WTP Bounds to Treatment of Insignificant 1st Stage Results

	(4)	(5)	(6)	(7)
model	IIV	IIV 95% CI	IIV 95% CI	IIV 95% CI
dependent variable in price function	price	price	price	ln(price)
demographics in demand function	all	all	income only	all
	clarity improves from 4.7 to 5.4 meters			
baseline: insignificant estimates for MWTP excluded	$2,679	$2,903	$2,944	$4,593
insignificant estimates for MWTP included	$2,470	$2,551	$2,566	$4,246
insignificant estimates for MWTP included; negative estimates set to zero	$2,477	$2,558	$2,572	$4,269
insignificant estimates for MWTP included; all set to zero	$2,447	$2,529	$2,541	$4,249
	clarity improves from 2.1 to 5.4 meters			
baseline: insignificant estimates for MWTP excluded	$26,426	$29,719	$29,157	$67,154
insignificant estimates for MWTP included	$25,547	$28,762	$28,834	$63,909
insignificant estimates for MWTP included; negative estimates set to zero	$25,561	$28,772	$28,845	$63,938
insignificant estimates for MWTP included; all set to zero	$25,601	$28,833	$28,948	$63,996

Note: The numbers in the table are upper bounds on the range of consistent IIV estimates for the average consumer's surplus for an improvement in water clarity. In each section of the table, the first row replicates our baseline results from table 5 in which the four markets with insignificant first-stage estimates for MWTP from table 4 are excluded. The second row includes data from the markets with insignificant estimates. The third row includes insignificant estimates with the restriction that all negative and insignificant estimates for MWTP are set to zero. The last row includes insignificant estimates but sets all insignificant estimates for MWTP to zero.