Using Panel Data To Easily Estimate Hedonic Demand Functions[*]

Kelly C. Bishop[†]
Arizona State University

Christopher Timmins[‡]
Duke University and NBER

February 2016

Abstract

The hedonics literature has often asserted that if one were able to observe the same individual make multiple purchase decisions, one could recover rich estimates of preferences for a given amenity. In particular, in the face of a changing price schedule, observing each individual twice is sufficient to recover the underlying (linear) demand function separately for each individual, with no restrictions on this heterogeneity in either the intercept or the slope. Using a rich panel dataset, we recover the full distribution of demand functions for clean air in the Bay Area of California. First, we find that estimating the full demand function, rather than simply recovering an estimate of local marginal willingness to pay, is important. Second, we find evidence of considerable heterogeneity in these functions and find that this heterogeneity is important from a policy perspective – our data-driven estimates of the welfare effects associated with a non-marginal change in air quality differ substantially from the existing approaches to welfare estimation.

Key Words: Hedonic Demand, Valuation, Ozone

JEL Classification Numbers: Q50, Q51, R21, R23

[*]We thank seminar participants at the Triangle Resource and Environmental Economics Workshop, Camp Resources, University of Wisconsin, Iowa State University, AERE Winter Meetings, Syracuse University, Montana State University, Washington University in St. Louis, University of Missouri St. Louis, and St. Louis University for their helpful comments and suggestions.

[†]Kelly Bishop, Arizona State University, Tempe, AZ 85287. kelly.bishop@asu.edu

[‡]Chris Timmins, Duke University, Box 90097, Durham, NC 27708. christopher.timmins@duke.edu

1 Introduction

Many applications of the hedonic model seek to value non-marginal changes in amenities, requiring the estimation of the underlying hedonic demand or marginal willingness-to-pay (MWTP) function. This, however, is not without costs as assumptions are typically needed to restrict preference heterogeneity. In this paper, we show how to recover the unconditional distribution of linear MWTP functions with a simple and transparent data-driven estimation approach.

The traditional estimation approach for recovering hedonic demand is based on Rosen's seminal 1974 paper. However, the second stage of Rosen's two-step approach suffers from a number of econometric problems; of particular concern is when the hedonic price function is non-linear in the amenity of interest, buyers simultaneously choose both the hedonic price and the quantity of that amenity that they will consume. Therefore, the problem of consumer choice subject to a non-linear budget constraint creates a difficult endogeneity problem when using statistical inference to recover the parameters describing those preferences.[1] Typically, IV approaches in this literature have relied upon questionable exclusion restrictions, e.g., certain socio-demographic variables enter directly into the MWTP function while others do not and the excluded variables can be used as instruments for endogenous attribute levels.[2] These assumptions are not testable and place arbitrary restrictions on the estimated heterogeneity in MWTP and, aside from market dummies, these instruments are generally hard to justify. These difficulties have led most researchers to forgo estimating the MWTP function altogether, employing only local measures of MWTP in policy analysis.

In their 2005 paper, Bajari and Benkard demonstrate that this endogeneity problem may be avoided by replacing the statistical inference used in the second stage of Rosen's method with a "preference inversion" procedure that inverts the first-order conditions of

[1]See Brown and Rosen (1982), Mendelsohn (1985), Bartik (1987), and Epple (1987) for discussions about issues relating to the identification and estimation of the MWTP function. More recently, Ekeland, Heckman, and Nesheim (2004) clearly demonstrate the conditions under which the MWTP function is identified. Bishop and Timmins (2015) build upon the insights of Ekeland, Heckman, and Nesheim to illustrate how estimating the MWTP function can be relatively straightforward. However, in both cases, restrictions on the heterogeneity of preferences are required. Banzhaf (2015) shows that by assuming a single-crossing condition, a first-order approximation to the MWTP function can be obtained.

[2]Alternatively, the MWTP function may only contain socio-demographic variables in linear form, while higher-order terms are excluded from the function and can therefore serve as instruments.

utility maximization to recover demand at the individual level. The strengths of this approach lie in (i) its admission of any form of preference heterogeneity and (ii) its avoidance of the endogeneity problems described above. Its primary weakness, however, comes in the strict functional-form assumptions that are required to perform the inversion procedure, i.e., when observing each individual on only one purchase occasion (one data point), it is only possible to recover the MWTP function by fully assuming its shape. When the shape of the MWTP function is so strongly dictated by functional-form assumptions, the value in going beyond the first stage of Rosen's two-step procedure is limited. Bajari and Benkard recognize this, pointing-out that these assumptions could be relaxed if the researcher were able to observe the same individual buyer on multiple purchase occasions.

Figure 1: Recovering Hedonic Demand Functions using Panel Data

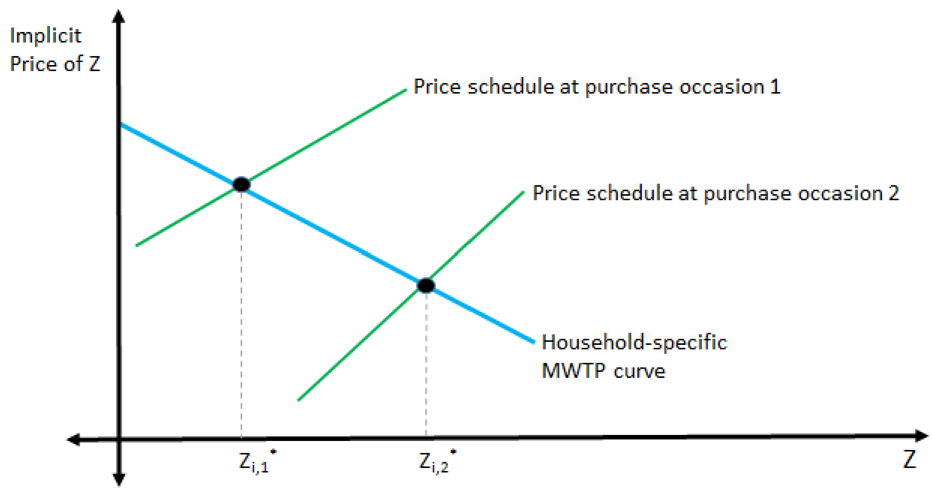

Our method for recovering the MWTP function is quite intuitive: observing the exact same household make purchase decisions in multiple geographic markets or in multiple time periods (i.e., under different price schedules), allows us to trace-out household-specific hedonic demand functions. Given a graphical representation of the problem, our empirical approach becomes clear; observing each household on at least two purchase occasions allows us to "connect the dots" and recover a unique (linear) MWTP curve for each household in the dataset. Figure (1) describes this approach for a given housing amenity, Z; when facing the associated (implicit) price schedules on purchase occasions 1 and 2, household

i chooses to consume $Z^*_{i,1}$ and $Z^*_{i,2}$, respectively. These two observations are sufficient for recovering a household-specific linear MWTP function for household i.

Importantly, the intuition behind our identification is straightforward. The schedule of market prices (i.e., the hedonic price gradient) exhibits exogenous variation across time from the point of view of any given household, as we assume that households are price-takers in the real estate market. Thus, we are observing each individual household choosing how much of amenity Z to consume under different supply conditions, allowing us to trace out the MWTP curve for each household in the data. Naturally, the MWTP function that we recover with two observed purchases is linear. However, we impose no additional restrictions on the shape of MWTP and, in particular, no restrictions on the heterogeneity of the individual-specific coefficients for intercept and slope. We show how these preference coefficients may be decomposed to isolate the effect of fixed demographic characteristics, such as race. We also show that with three or more observations for each household, we can estimate the effect of time-varying demographic characteristics, such as income.[3]

We apply this methodology to recovering individual-specific MWTP functions for air quality in the Bay Area of California. In particular, we estimate the marginal willingness to pay to avoid ground-level ozone pollution. For this analysis, we create a rich panel dataset describing real estate transactions and the attributes of the associated buyers over the thirteen-year period 1991-2003. To create this dataset, we first isolate and match individual households over time and then merge the demographic characteristics provided on mortgage applications per the Home Mortgage Disclosure Act of 1975. Finally, we generate a house-level annual measure of ozone pollution from the monitor data provided by the California Air Resources Board. In particular, our measure of ozone is days exceeding the California state maximum 1-hour ozone concentration.

In the implementation of this approach, we allow for the most flexible representation of preferences possible with the available data. We begin by estimating a non-parametric regression to recover a flexible set of time-varying hedonic price gradients which includes a full set of fixed-effects at the Census-tract level. In the second stage, we estimate household-specific, linear MWTP curves.

As we estimate MWTP functions (versus the local measure only), we can calculate the elasticity of the MWTP with respect to ozone exposure. We find that the median

[3]With three or more observations per household, we could alternatively allow for greater flexibility in the shape of the MWTP curve, such as allowing the function to be quadratic in Z.

value of this elasticity is 0.68, implying that it is important to recover the entire function when considering non-marginal changes. We also find considerable heterogeneity across households in the MWTP to pay to avoid ozone pollution with an interquartile range of 0.85 for the previously-described elasticity.

Correctly estimating heterogeneity in the slope of the MWTP function is important from a policy perspective as the policy-maker must account for the fact that those exposed to larger policy-induced changes in amenities may be more (or less) sensitive to the change, compared with the average household. In other words, do the households exposed to the largest changes have the steeper or flatter MWTP functions? This question is particularly important given that households likely sort based on their preferences for ozone.

To illustrate this concept, we estimate the welfare costs of a non-marginal increase in ozone. We find that the mean willingness to pay to avoid a 33% increase in ozone is \$953.80 and that the interquartile range is \$1,064.25. We first compare our results with a specification that allows for heterogeneity in MWTP intercepts, but eliminates heterogeneity in MWTP slopes by assuming all households have flat MWTP functions. Using these flat MWTP functions returns the expected result of smaller mean welfare costs. In this case, the average cost is \$796.30. We then compare our results with a specification where the slope is identified by the functional form of the utility function, as in Bajari and Benkard (2005). In this case, both the estimated average welfare costs and the estimated variance of welfare costs are considerably larger: the mean is \$2,437.22 and the interquartile range is \$2,072.86. This difference is explained by the negative correlation between change in ozone exposure and the slope of the MWTP function. The functional-form based approach recovers a negative correlation; the steepest MWTP curves are assigned to those households which are exposed to the highest amount of ozone. In contrast, when we use panel data and households' observed changes in ozone exposure (in response to observed changes in price) to recover the slope of the MWTP curve, we obtain a positive correlation between change in ozone exposure and the slope of the MWTP function.

This paper proceeds as follows. Section 2 describes our methodological approach for recovering the hedonic demand for air quality in the Bay Area of California. The creation of our unique two-sided panel dataset and its summary statistics are discussed in Section 3. Section 4 presents our results and Section 5 applies our MWTP-function results to measure the willingness to pay to avoid a non-marginal change in ozone levels. Finally, Section 6 concludes.

2 Model

In this section, we describe our panel-data-driven approach to recovering the structural parameters of the hedonic model. First, we discuss the non-parametric repeat-sales model that we use to recover the hedonic gradient with respect to a particular amenity. While this approach allows for the fewest assumptions and still controls for house-level fixed-effects, a simpler specification would also be sufficient for the identification of the second stage. We then show how with access to a panel of buyers (and implicit prices from the first stage), we can recover fully heterogeneous MWTP functions by observing households in at least two different time periods.

2.1 A Non-Parametric Fixed-Effects Approach to Recovering the Hedonic Gradient

In the first stage of the model, we estimate the hedonic price function, which relates the price of house j transacted in period t ($P_{j,t}$) to its attributes: both those that vary over time, $Z_{j,t}$, and those that do not, X_j. The gradient may be estimated in any number of ways, although a simple, linear framework may impose unrealistic restrictions on the equilibrium underlying the hedonic price function.[4] Additionally, concern must be paid to the potential bias caused by omitted variables.[5] Using panel data with repeat sales and controlling for house fixed effects avoids the bias caused by time-invariant house characteristics, whether observed or unobserved.

Following Fan and Gijbels (1996) and Lee and Mukherjee (2014), we adopt a non-parametric first-difference approach for the estimation of the gradient and begin by writing down a flexible representation of the hedonic price function:

$$P_{j,t} = f(Z_{j,t}) + X_j + \nu_{j,t} \tag{1}$$

where $f(\cdot)$ is an unspecified, flexible function of time-varying attributes of house j (or its

[4]See Heckman, Ekeland, and Nesheim (2004) for a discussion of the shape of equilibrium hedonic price functions.

[5]Bias from omitted variables arises when the regressors are not orthogonal to the regression error. This would be the case if data describing important neighborhood or housing attributes (*e.g.*, distance to city center or curb appeal) are not available to the researcher, but those variables are correlated with the attribute of interest.

neighborhood) and X_j represents all time invariant attributes (whether they are observed by the econometrician or not). While we allow $Z_{j,t}$ to be correlated with the (potentially unobservable) elements of X_j, we assume that $Z_{j,t}$ is uncorrelated with the time-varying error term, $\nu_{j,t}$. We allow $\nu_{j,t}$ to be correlated with X_j. In our application, $Z_{j,t}$ will consist of (i) a measure of ground-level ozone pollution at house j in year t and (ii) the year of the housing transaction.[6]

We take a first-order Taylor series expansion of $f(\cdot)$ around a vector χ, which has the same dimension as $Z_{j,t}$.[7]

$$P_{j,t} = f(\chi) + (Z_{j,t} - \chi)' f'(\chi) + X_j + \nu_{j,t} \tag{2}$$

We denote each property's prior year of sale by t_{-1} and rewrite Equation (2) for this prior sale:[8]

$$P_{j,t_{-1}} = f(\chi) + (Z_{j,t_{-1}} - \chi)' f'(\chi) + X_j + \nu_{j,t_{-1}} \tag{3}$$

This allows us to subtract Equation (3) from Equation (2) and difference-out both the time-invariant attributes X_j and the time-invariant term $f(\chi)$.[9]

$$P_{j,t} - P_{j,t_{-1}} = (Z_{j,t} - Z_{j,t_{-1}})' f'(\chi) + (\nu_{j,t} - \nu_{j,t_{-1}}) \tag{4}$$

Denoting first differences with "$\sim$" and replacing $f'(\cdot)$ with $\beta(\cdot)$, we arrive at:

$$\tilde{P}_{j,t} = \tilde{Z}'_{j,t} \beta(\chi) + \tilde{\nu}_{j,t} \tag{5}$$

Lee and Mukherjee (2014) show that $\beta(\chi)$ (i.e., the slope of the hedonic price func-

[6]In practice, we use the number of days in the course of each year that the state maximum 1-hour ozone concentration of 90 ppb is violated. Other potential measures of ground-level ozone include the 1-hour and 8-hour average maximum concentrations over the course of each year. Ozone exceedances are a better measure of extreme pollution events that house buyers may be more aware of.

[7]The remainder term associated with the Taylor expansion is ignored.

[8]This is expanded around the same vector χ.

[9]Losing $f(\cdot)$ from this expression does not pose a problem, as our interest is only in recovering the hedonic gradient, *i.e.*, the slope of the hedonic price function, which is represented non-parametrically by $f'(\cdot)$.

tion at $Z_{j,t} = \chi$) may be recovered with the following least-squares minimization procedure:

$$\beta(\chi) = arg \min_{\beta(\chi)} \sum_{j=1}^{J} \sum_{t=1}^{T_j} \left(\tilde{P}_{j,t} - \tilde{Z}'_{j,t}\beta(\chi) \right)^2 K_h(Z_{j,t} - \chi) K_h(Z_{j,t_{-1}} - \chi) \tag{6}$$

where T_j denotes the number of first-differenced observations for each house j. Implementing the local-linear regression procedure requires choosing the weights placed on data as one moves further from the point of evaluation, χ. For our estimation, $K_h(\cdot)$ is given by the Gaussian kernel:

$$K_h(Z_{j,t} - \chi) = \prod_{k=1}^{2} \frac{1}{h\hat{\sigma}_{Z_k}} \frac{1}{\sqrt{2\pi}} exp\left\{ -\frac{1}{2} \left(\frac{Z_{k,j,t} - \chi_k}{h\hat{\sigma}_{Z_k}} \right)^2 \right\} \tag{7}$$

where h represents the kernel bandwidth and $\hat{\sigma}_{Z_k}$ is the standard deviation of the k^{th} element of $Z_{j,t}$.

In practice, this allows us recover an estimate of $\beta(\chi)$ for all observed values of $Z_{j,t}$. We therefore end up with a (potentially) different estimate of $\beta(\chi)$ at each data point; this is in contrast to a fully parametric estimation procedure, where β would be constrained to be the same for every value of $Z_{j,t}$.

Given the linear nature of this problem, estimation of $\beta(\chi)$ can be summarized as the following weighted least-squares regression:

$$\beta(\chi) = (\tilde{Z}' W_\chi \tilde{Z})^{-1} \tilde{Z}' W_\chi \tilde{P} \tag{8}$$

where $n = \sum_{j=1}^{J} T_j$ is the total number of first-differenced observations, $\tilde{Z}$ is an (n x 2) matrix of first-differenced (within-house) regressors, $\tilde{P}$ is an (n x 1) vector of first-differenced (within-house) house prices, and W_χ is an (n x n) matrix of weights, such that $W_\chi = diag\big(K_h(Z_{j,t} - \chi)\big)$.

2.2 Recovering MWTP Functions using a Panel of Buyers

In this section, we demonstrate that it is straightforward to recover the flexible distribution of linear MWTP functions with panel data on home buyers. We begin by specifying the utility of household i with non-housing consumption ($C_{i,t}$), choosing a house j with attributes $(X_j, Z_{j,t})$, in period t:

$$U(X_j, Z_{j,t}, C_i) = \alpha_{0,i} + \alpha_{1,i}X_j + \frac{\alpha_{2,i}}{2}X_j^2 + \alpha_{3,i}Z_{j,t} + \frac{\alpha_{4,i}}{2}Z_{j,t}^2 + C_{i,t} \tag{9}$$

where we have normalized the coefficient on consumption to one.

Household i faces a budget constraint, $C_{i,t} + R_{j,t} \leq I_{i,t}$, where $I_{i,t}$ denotes income and $R_{j,t}$ is the imputed annual rent or housing expenditure associated with house j. In practice, we calculate this figure as 5% of the observed transaction price.[10]

Maximizing utility dictates that household i's budget constraint will bind, so we can incorporate the budget constraint $(C_{i,t} + R_{j,t} = I_{i,t})$ to arrive at the indirect utility function:

$$V_{i,j,t} = \alpha_{0,i} + \alpha_{1,i}X_j + \frac{\alpha_{2,i}}{2}X_j^2 + \alpha_{3,i}Z_{j,t} + \frac{\alpha_{4,i}}{2}Z_{j,t}^2 + (I_{i,t} - R_{j,t}) \tag{10}$$

We now demonstrate that this functional form will yield the same linear MWTP specification that is common in the hedonics literature. Moreover, as long as MWTP is not a function of time-varying household attributes, the parameters of the MWTP function can be identified with just two observations for each household. The first-order condition associated with the household's optimal choice of $Z_{j,t}$ is given by:

$$\frac{\partial V_{i,j,t}}{\partial Z_{j,t}} = \alpha_{3,i} + \alpha_{4,i}Z_{j,t} - \frac{\partial R_{j,t}}{\partial Z_{j,t}} = 0 \tag{11}$$

and household i's MWTP function for $Z_{j,t}$ is given by:

$$MWTP(Z_{j,t}) = \alpha_{3,i} + \alpha_{4,i}Z_{j,t} \tag{12}$$

For household i, we need to recover the values of two unknown parameters: $(\alpha_{3,i}, \alpha_{4,i})$, or the intercept and slope of MWTP, respectively.[11] Fortunately, we observe household i in panel data on (at least) two occasions. For households observed twice, we then have two

[10]This is a commonly used discount in the literature. See Poterba (1984) for a discussion of converting prices to annualized user-cost measures.

[11]Murray (1975) takes the approach of using an equation very similar to (11), except without individual heterogeneity in preference parameters, to form an estimating equation. Recovering estimates of $\frac{\partial R_{j,t}}{\partial Z_{j,t}}$ from the first-stage, he then recovers estimates of α_3 and α_4 from a two-stage least squares procedure using income and prices as instruments.

equations in two unknowns:

$$\alpha_{3,i} + \alpha_{4,i} Z_{j^*(i),1} = \rho_{j^*(i),1} \qquad \alpha_{3,i} + \alpha_{4,i} Z_{j^*(i),2} = \rho_{j^*(i),2} \tag{13}$$

where $\rho_{j^*(i),t} = \left.\frac{\partial R_{j,t}}{\partial Z_{j,t}}\right|_{Z_{j,t}=Z_{j^*(i),t}}$. Estimates of $\rho_{j^*(i),t}$ are recovered in the first-stage non-parametric regressions and $t = 1, 2$ are the two periods in which household i purchases.

Solving these two equations yields closed-form solutions for the structural parameters of the utility function, i.e., the household-specific solutions for both the intercept and the slope of the MWTP function. Respectively, these are given by:

$$\hat{\alpha}_{3,i} = \frac{\hat{\rho}_{j^*(i),2} Z_{j^*(i),1} - \hat{\rho}_{j^*(i),1} Z_{j^*(i),2}}{Z_{j^*(i),1} - Z_{j^*(i),2}} \qquad \hat{\alpha}_{4,i} = \frac{\hat{\rho}_{j^*(i),1} - \hat{\rho}_{j^*(i),2}}{Z_{j^*(i),1} - Z_{j^*(i),2}} \tag{14}$$

2.2.1 Allowing MWTP to Vary with Individual Household Attributes

It is possible to determine how these MWTP functions differ systematically with fixed individual attributes, A_i (e.g., race, completed education, or gender).[12] This is easily done by performing the following least-squares regressions in a separate stage:

$$\hat{\alpha}_{3,i} = \delta_{3,0} + A_i' \delta_{3,1} + \varsigma_{3,i} \qquad \hat{\alpha}_{4,i} = \delta_{4,0} + A_i' \delta_{4,1} + \varsigma_{4,i} \tag{15}$$

However, income is often included as a determinant of MWTP, but is usually time-varying (in panel data) and hence cannot be summarized by the vector of fixed attributes, A_i. Murray (1983) shows why it is important for the MWTP function to vary with non-housing consumption expenditure. We demonstrate next how time-varying, non-housing consumption expenditure can be included with access to *three* observations on each individual household. We begin by specifying household i's indirect utility from choosing

[12]If A were time varying, we would need to observe the household on more purchase occasions. For example, if A_i were a time-varying scalar that affected both the intercept and slope of MWTP, preference heterogeneity could be represented as $\alpha_{3,i,t} = \gamma_{3,0,i} + \gamma_{3,1,i} A_{i,t}$ and $\alpha_{4,i,t} = \gamma_{4,0,i} + \gamma_{4,1,i} A_{i,t}$. A household would need to be observed on four occasions, with Equation 11 yielding four equations which could be solved for the four unknowns, $\gamma_{3,0,i}$, $\gamma_{3,1,i}$, $\gamma_{4,0,i}$, $\gamma_{4,0,i}$. If $A_{i,t}$ were to shift only the slope or the intercept of the MWTP function (but not both), the household would only need to be observed on three purchase occasions.

house j in period t as:

$$V_{i,j,t} = \alpha_{0,i} + \alpha_{1,i} X_j + \frac{\alpha_{2,i}}{2} X_j^2 + \alpha_{3,i} Z_{j,t} + \frac{\alpha_{4,i}}{2} Z_{j,t}^2 + \alpha_{5,i} Z_{j,t}(I_{i,t} - R_{j,t}) + (I_{i,t} - R_{j,t}) \quad (16)$$

where we now allow the marginal utility of $Z_{j,t}$ to vary with non-housing consumption expenditure.[13] The first-order condition associated with the individual's optimal choice of $Z_{j,t}$ is then given by:

$$\frac{\partial V_{i,j,t}}{\partial Z_{j,t}} = \alpha_{3,i} + \alpha_{4,i} Z_{j,t} + \alpha_{5,i}(I_{i,t} - R_{j,t}) \underbrace{-\alpha_{5,i} Z_{j,t} \frac{\partial R_{j,t}}{\partial Z_{j,t}} - \frac{\partial R_{j,t}}{\partial Z_{j,t}}}_{-\rho_{j,t}(\alpha_{5,i} Z_{j,t}+1)} = 0 \quad (17)$$

and household i's MWTP function for $Z_{j,t}$ is given by:

$$MWTP(Z_{j,t}) = \frac{\alpha_{3,i} + \alpha_{4,i} Z_{j,t} + \alpha_{5,i}(I_{i,t} - R(Z_{j,t}))}{(\alpha_{5,i} Z_{j,t} + 1)} \quad (18)$$

As $\alpha_{5,i}$ is presumably a very small number, this closely approximates a linear function of $Z_{j,t}$.

For households that are observed on three separate purchase occasions, Equation 17 holds for each of the three purchase occasions, yielding three equations in three unknowns: $\alpha_{3,i}$, $\alpha_{4,i}$, and $\alpha_{5,i}$. Appendix A.2 outlines the straightforward algebraic recovery of $\hat{\alpha}_{3,i}$, $\hat{\alpha}_{4,i}$, and $\hat{\alpha}_{5,i}$.

3 Data

To implement our first-stage nonparametric regressions, we use data describing single-family housing transactions over the period 1991 to 2003 in the Bay Area of California. For the second-stage recovery of the MWTP functions, we require data on a panel of home buyers. For this, we assemble a dataset by combining information from the real estate transactions dataset and a dataset describing mortgage applicants' demographic characteristics obtained through the Home Mortgage Disclosure Act (HMDA).

[13]Note that our indirect utility specification could be expanded to include interactions with multiple time-varying attributes. Were we to include these, we would simply need to solve for multiple first-order conditions (using multiple derivatives of the price function) for each observed purchase. In this paper, we keep things relatively simple and avoid this complication.

3.1 Property Transactions Data

The real estate transactions data we employ cover the six core counties of the San Francisco Bay Area (Alameda, Contra Costa, Marin, San Francisco, San Mateo, Santa Clara). The data were purchased from DataQuick and include transaction dates, prices, loan amounts, and buyers', sellers' and lenders' names for all transactions. In addition, the data for the final observed transaction include housing characteristics, such as exact street address, square footage, year built, lot size, number of rooms, number of bathrooms, and number bedrooms.

Table 1: Housing Transactions Data

	Full Sample		Repeat Sales Sample	
	n = 630,384		*n = 277,011*	
variable	mean	st.dev.	mean	st.dev.
Price (in year 2000 $)	379,368.70	202,887.70	371,189.30	192,626.00
Sq. Ft. House	1,716.42	679.78	1,651.01	634.30
Sq. Ft. Lot	7,148.83	7,962.83	6,523.88	7,068.32
Year Built	1966.92	22.66	1968.19	21.67
Num. Bedrooms	3.20	0.90	3.12	0.88
Num. Bath	2.12	0.73	2.09	0.69

As housing characteristics are only provided for the final assessment of each property, we take steps to ensure that the house has not undergone any major changes. First, to control for land sales or re-builds, we drop all transactions where "year built" is missing or with a transaction date that is prior to "year built." Second, to control for major property improvements that would not present as a re-build, we drop properties that experience a yearly appreciation/depreciation rate that is more than four times greater than the average appreciation/depreciation rate (in absolute value).[14] Additionally, we drop transactions where the price is missing, negative, or zero. After using the consumer price index to convert all transaction prices into 2000 dollars, we drop one percent of observations from each tail to minimize the effect of outliers. As we merge in the pollution data using the property's geographic coordinates, we drop properties where latitude and longitude are

[14] In a repeat-sales analysis, similar in spirit to Case and Shiller (1989), we regress log prices on a set of house and year dummies, which provides us a crude measure of yearly appreciation (or depreciation) rates in the Bay Area.

missing.[15]

Finally, we restrict our analysis to properties with multiple sales over the thirteen year period.[16] This yields a final sample of 277,011 transactions (*i.e.*, property-year observations) comprised of 126, 227 unique properties. Table 1 describes the data.

3.2 Buyer Characteristics Data

In order to implement our second-stage estimator, we need to create a panel of households, their characteristics, and their chosen properties over time. This involves first identifying and matching households over time in the property transactions dataset and, second, merging-in the attribute data from the HMDA dataset. This is possible as we have buyer and seller names in the transaction record as well as common variables in both datasets.

To track households over time we use the buyer's name. To minimize the probability of finding erroneous matches, we only accept a match if the entire name field provided in the dataset matches over time. In most cases, the entire name includes a first and last name for both the primary buyer and the secondary buyer (i.e., usually the spouse) and it often includes middle initials.[17] As counties may record data differently, we also only use matches where purchases occurred within the same county.[18] Finally, we keep only observations where the buyer is observed on either two or three purchase occasions.[19]

[15]Although we use a repeat-sales estimation approach, we make some cuts on property characteristics to ensure all single-family homes are trading in the same market: dropping properties where year built is pre-1850, lotsize is either zero or greater than three acres, square footage is either less than 400 or greater than 10,000, number of bedrooms or bathrooms is greater than ten, number of total rooms is greater than fifteen, or number of stories is greater than three.

[16]However, we drop properties that sell more than once within a calendar year or more than five times over the thirteen year period.

[17]As a robustness check, we also consider a weaker version of the algorithm where we require a match only on (the broken-out variables of) first and last names and the results are similar. As discussed in the Results section, with this less stringent matching algorithm, there is a small increase in the number of second-order-condition violations, which is suggestive that using the most stringent algorithm is appropriate.

[18]Again, the results are robust to this rule. However, as with the name merge, when we use the less-stringent matching algorithm of allowing across-county matches, there is a small increase in the number of second-order-condition violations. Again, this is suggestive that only using within-county moves is appropriate.

[19]If a household purchases less than twice in our sample period, we cannot use the information to pin down the parameters of their MWTP function. If we observe a buyer on more than three occasions (over a thirteen-year sample period), we assume he/she is a professional and drop them from the sample. Similarly, we drop buyers who are observed buying more than once in a calendar year.

Table 2: One-, Two-, and Three- Purchase Samples

	One-Purchase *households=315,215*		Two-Purchase *households=6,073*		Three-Purchase *households=220*	
variable	mean	median	mean	median	mean	median
Asian	0.24	0	0.26	0	0.29	0
Black	0.03	0	0.03	0	0.01	0
Hispanic	0.11	0	0.13	0	0.15	0
White	0.61	1	0.58	1	0.55	1
Income (in yr 2000 $)	114,464	96,014	125,709	103,370	132,761	109,224
Price (in yr 2000 $)	390,801	347,580	413,738	359,238	394,637	329,543

The individual attribute data come from a dataset on mortgage applications published through the Home Mortgage Disclosure Act of 1975. These data provide information on all mortgage applications filed in the Bay Area over the period of our sample. Included are all applicants' race and gender, income, loan amount, lender name, and Census tract of the property. We are able to merge the individual attributes in the HMDA dataset to the buyers in the property transactions dataset using the common variables of lender name, loan amount, transaction date, and Census tract of the property.[20] We successfully match approximately two-thirds of households in the transactions sample to the raw HMDA sample.[21] Table 2 provides statistics describing the one-purchase, two-purchase, and three-purchase buyers.

3.3 Ozone Data

The ozone data we employ are taken from the California Air Resources Board.[22] We use yearly ozone data from thirty-five monitors in the nine counties of Alameda, Contra Costa, Marin, Napa, San Francisco, San Mateo, Santa Clara, Solano, and Sonoma over the period 1990 to 2004. In particular, we use the monitor data to construct property-specific

[20]See Bayer, McMillan, Murphy, and Timmins (forthcoming) for more details on this merge.

[21]After the merge, we drop observations where either race or income is missing. As the race variable is self-reported at each purchase, it is possible that a buyer reports a different race in different periods. In addition to buyers simply misreporting their own race, this could be driven by issues with the DataQuick-HMDA merge or with the buyer-tracking algorithm. However, these cases are not prevalent in the data. Buyers report the same race in 92% of cases. For over half of the buyers who report a change in race, the change in either to or from Hispanic (which is a race category in the data). The results presented in the next section are robust to excluding those who report a change in race.

[22]Publicly available at www.arb.ca.gov/adam/.

measures of the number of days exceeding the one-hour California standard (*i.e.*, 90 parts per billion).

In addition to the ozone readings, the dataset provides information on the "year coverage," or the percent of time (during the relevant high-ozone season) each particular monitor was available and the geographic coordinates of each monitor.[23] Using this coverage variable, we drop monitors with less than 60 percent coverage in a given year (amounting to less than 4 percent of the available monitor-year observations).

Using the latitudinal and longitudinal coordinates of the monitors and the properties, we use the "Great Circle" estimator to compute the distance to all monitors from each property. We then create a three-year weighted average for each property of all monitors' readings, weighting distance by one over distance-squared.[24] In order to mitigate boundary effects, we include monitor data from the surrounding counties of Napa, Solano, and Sonoma, in addition to the six counties that appear in our transactions data.

The maps in Figure 2 describe the spatial distribution of ground-level ozone pollution in the Bay Area. Two important features emerge from these pictures. First, geography is largely responsible for cross-sectional variation in pollution. San Francisco (on the tip of the peninsula extending from the South Bay into the Pacific Ocean), Oakland (in the East Bay), and San Jose (at the southern end of the San Francisco Bay) all face heavy traffic congestion. Wind patterns, however, mitigate much of the ozone pollution in San Francisco and Oakland, while worsening it in San Jose. Mountains ringing the southern end of the Bay Area block air flows and contribute to this effect. The mountains on the eastern side of the Bay are similarly responsible for high levels of pollution along the I-680 corridor in eastern Contra Costa and Alameda counties.

These maps also make clear that there is significant variation in pollution levels over time. Figure 3 describes this time variation. Much of this is due to a variety of programs that were initiated after California passed its Clean Air Act in 1988. After multiple years of relatively low ozone pollution, while the Bay Area counties were designated as being out of attainment according to EPA rules, the Bay Area experienced the worst air quality since the mid-eighties in 1995, corresponding to the time at which it was placed back into

[23]Some monitors were opened or permanently closed during this time period.

[24]We use a three-year average of ozone as price and homeowner behavior more likely reflects this slightly longer average of ozone rather than the short-run annual measures. Results are robust to this choice of data smoothing. See Bishop and Murphy (2015) for a discussion of using time-averaged amenities in hedonic analyses.

Figure 2: Ozone Pollution in the Bay Area

attainment and mandatory ozone reduction programs were eliminated. In 1996, the Vehicle Buyback Program for cars manufactured in 1975 or before was implemented. This program, in addition to the Lawn Mower Buyback and the Clean Air Plan of 1997, presumably contributed to falling ozone levels. With even stronger mandatory ozone reduction policies after 1998 (mostly targeted at mobile sources), the remaining years of our sample returned to relatively low ozone levels. Also during the late 1990s, almost 100 emitting facilities were reviewed under the Title V Program Major Facility Review. There is no reason to expect that any of these programs would have had special economic consequences for housing

Figure 3: Average Number of Days Exceeding the CA 1-Hour Standard

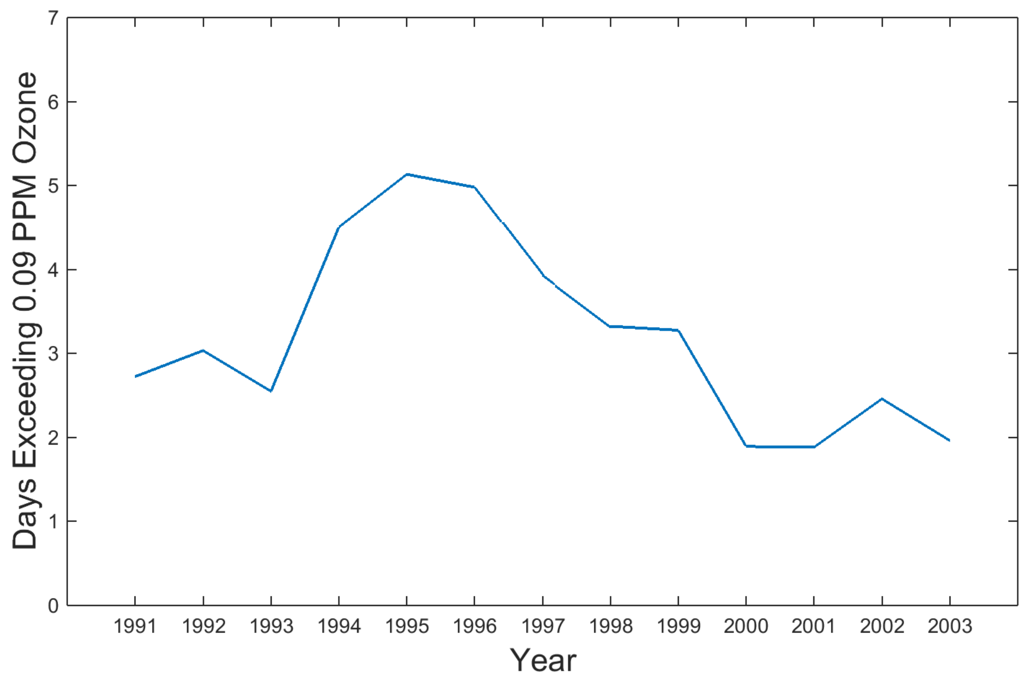

prices any particular part of the Bay Area, aside from those coming through changing amenity values.

4 Results

In this section, we describe our results. First, we illustrate the results of our non-parametric estimation procedure that is used to estimate hedonic gradients with respect to ground-level ozone pollution. Second, we recover the second-stage MWTP functions. In particular, we use the panel of 6,073 households that are observed purchasing two houses in our dataset to recover estimates of fully-heterogeneous MWTP functions as described in Section 2.2. We also use the sample of 220 households that buy three times to recover estimates of MWTP functions that are allowed to vary with non-housing expenditure. As a comparison, we show second-stage estimates of MWTP functions using two alternative approaches that do not require panel data and are commonly applied in the literature.

4.1 Results from the First-stage Hedonic Regressions

The local linear estimation allows the estimate of the slope coefficient, $\beta(\chi)$, to differ for each observed value of ozone pollution for each year of the panel. Thus, for each year, the gradient may be graphically represented as a flexible function of ozone pollution.

In estimation, we choose values of the bandwidths such that the gradients are monotonic.[25] While our results are robust to the choice of bandwidths, having monotonic gradients allows us to compare our results with approaches where the MWTP function is flat.[26]

Figure 4: Hedonic Price Gradients (1992, 1997, 2002)

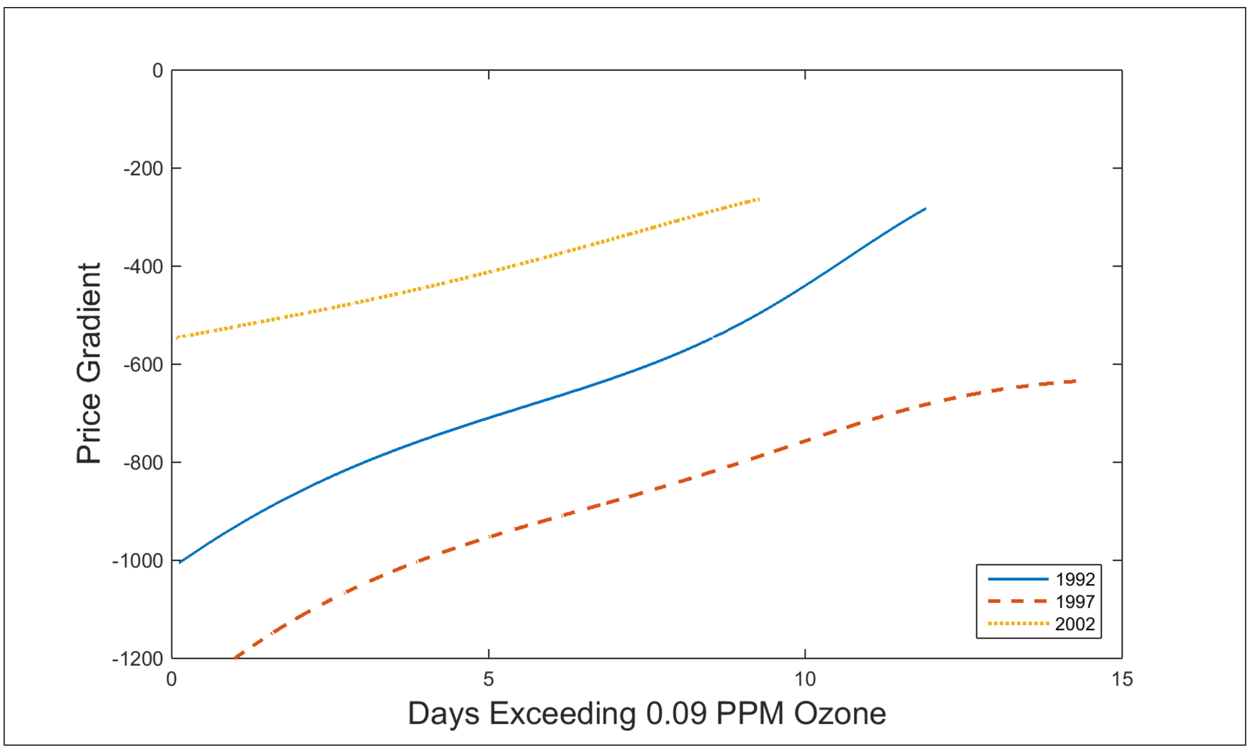

We estimate hedonic gradients for each of the thirteen years in our sample. For ease of exposition, we show the hedonic gradient for three of these years (1992, 1997, and 2002)

[25]This is achieved by setting the bandwidth over ozone to 4 and the bandwidth over year to 3.

[26]When the gradients are fully flexible, violations of the second-order conditions for utility maximization arise; households that choose a level of z which lies on a downward-sloping part of the gradient are not maximizing utility unless the MWTP function is sufficiently steep. In the case of the flat MWTP functions, this sufficient steepness cannot be achieved.

in Figure 4.1 and show all of the gradients with 95% confidence intervals in Figure A.1 in the Appendix.[27] The gradients are everywhere negative in sign (i.e., an increase in the level of pollution at a house leads to a reduction in its sale price, *ceteris paribus*). Moreover, these negative gradients are increasing in the level of pollution. This suggests that households may sort based on preferences (i.e., those with a lower willingness to pay tend to live in places with greater levels of pollution). Another feature of the data becomes evident when comparing Figure 4 with Figure 3: as expected, average exposure to ozone is higher in periods where the gradient is more negative. These features have important consequences for policy analysis and are discussed in detail in Section 5.

Finally, we show the density of gradients evaluated at the point of ozone exposure in Figure 5. Applying the first-order condition for utility maximization allows us to interpret these estimates as local measures of MWTP, holding locally at the point of observed exposure. The median value for the distribution presented in Figure 5 is -$696.29.

Figure 5: Distribution of Local Estimates of MWTP

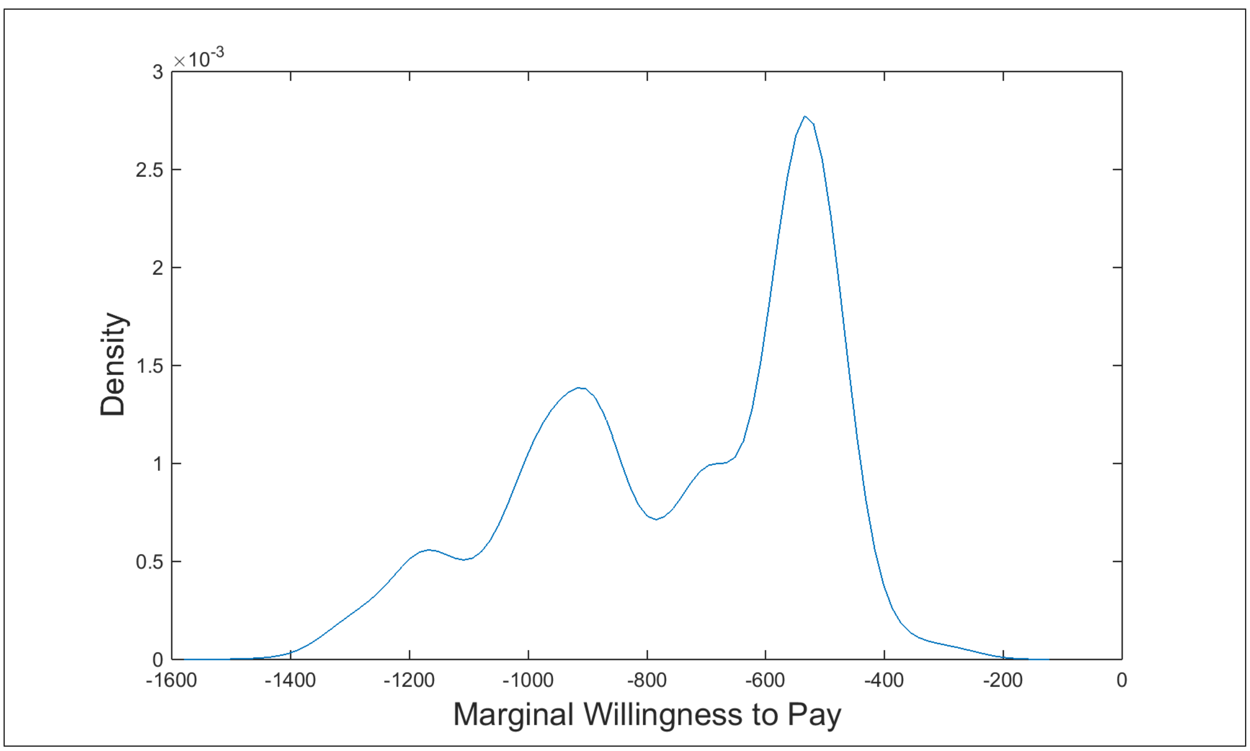

[27]Each gradient is plotted between the 1st through the 99th percentiles of ozone for that year.

4.2 Results from the Second-Stage MWTP Estimation

4.2.1 Demand Estimation with the Two-Purchase Panel

Employing our approach with a panel of households who are observed on two separate purchase occasions allows for the recovery of both an intercept and a slope coefficient for the MWTP function for each household. Thus, each household's MWTP function is allowed to vary with their level of ozone exposure, which can have important implications for valuing large changes in pollution. The fully heterogeneous slopes and intercepts are recovered according to the formulas in Equation 14.[28]

In a small minority of cases (20%), the recovered slopes are positive and sufficiently steep that they exceed the positive slope of the gradients, thus violating the second-order conditions for maximization. As these households' estimated behavior is not consistent with utility maximization under the assumptions of our model, we drop them from the analysis that follows. Two potential factors could cause us to incorrectly recover slopes which would explain why some households are not maximizing utility. The first is that preferences may not be stable over time and the second is that the merge on buyer name in the data is potentially imperfect. While our data is novel, even richer data could potentially address both these issues. For example, richer time-varying attribute data would allow our MWTP curve to vary over time based on observable characteristics. Additionally, a dataset that formally tracked households over time would allow us to avoid the name merge altogether; while we use the strictest algorithm for tracking households, we have found that less-strict algorithms lead to small increases in the number of second-order-condition violations, which is suggestive that some households are imperfectly tracked.

We find considerable evidence that MWTP varies with ozone with a median slope value of -\$191.95. However, it is difficult to summarize the joint distributions of intercepts and slopes in a single figure. We instead report the distribution of MWTP elasticities (with respect to ozone) in Figure 6. More specifically, this elasticity is calculated as $\hat{\alpha}_{4,i}\frac{Z_{j^*(i),t}}{\hat{\rho}_{j^*(i),t}}$.

The elasticity of MWTP with respect to ozone exposure is consistently positive, implying that the (negative) MWTP for pollution gets larger (in absolute value) as pol-

[28] As the second stage is an inversion, rather than an estimation, the only sampling variance of the estimates comes from the first stage, where the price function is very precisely estimated (see Figure A.1). As such, standard errors for the second-stage MWTP estimates are not reported.

Figure 6: Distribution of MWTP Elasticities

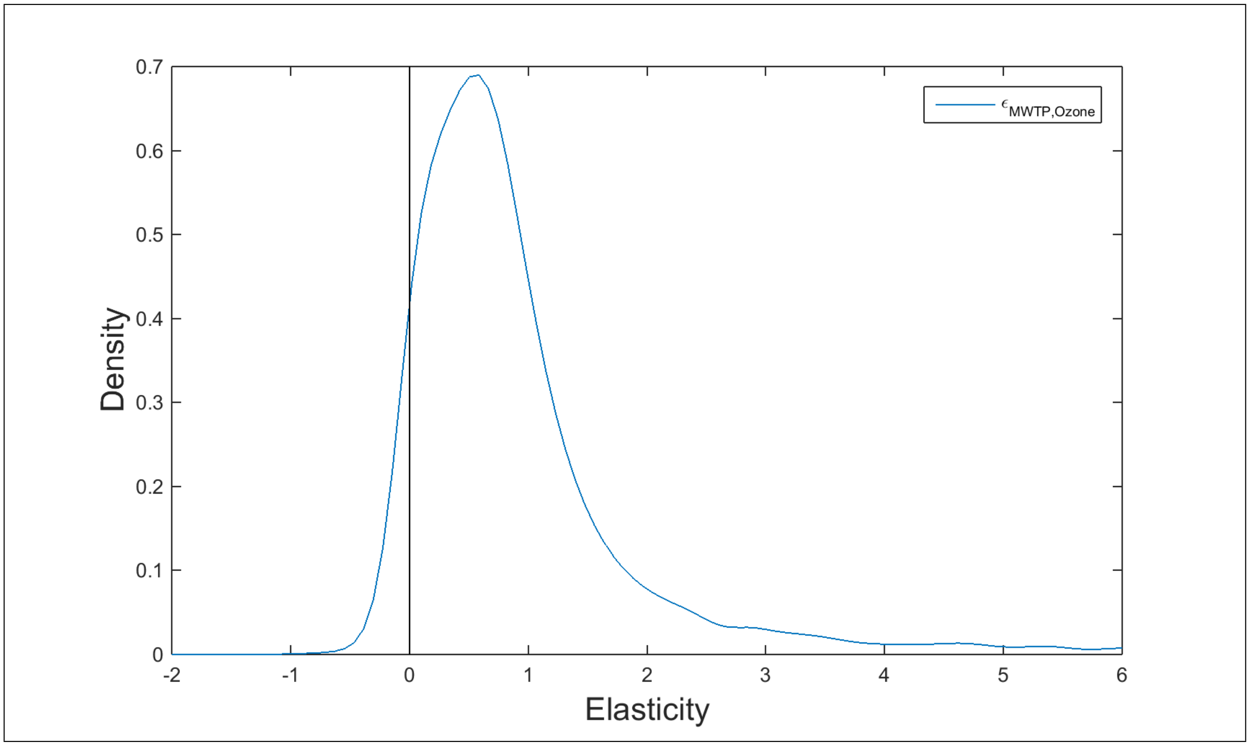

lution increases, with a median value of 0.68.[29] The distribution of these elasticities also shows significant heterogeneity. While the median value is 0.68, the distribution has an inter-quartile range of 0.85. This elasticity is not, however, strongly correlated with the level of ozone pollution people are experiencing (i.e., a correlation of 0.02).

Finally, we can decompose the intercept ($\hat{\alpha}_{3,i}$) and slope ($\hat{\alpha}_{4,i}$) coefficients according to Equation 15. In Figure 7, we plot the mean MWTP functions conditional on race for each of the observed races in our data: Asian, Black, Hispanic, and White.[30] From the figure, we can see that Asians and Whites have steeper MWTP functions compared with Blacks and Hispanics. This translates into a higher absolute value for MWTP only at larger levels of ozone exposure: for levels of ozone exposure below the median of 2.59, the MWTP is roughly similar for all groups.

[29]In a very small number of cases (5%), the elasticity is negative. This corresponds to an upward sloping MWTP function that still satisfies the second-order conditions for utility maximization. That is, while both marginal utility and the price gradient are increasing, marginal utility is increasing at a slower rate and so an interior solution exists.

[30]Each MWTP function is plotted between the 1st through the 99th percentiles of ozone exposure.

Figure 7: Average MWTP by Race

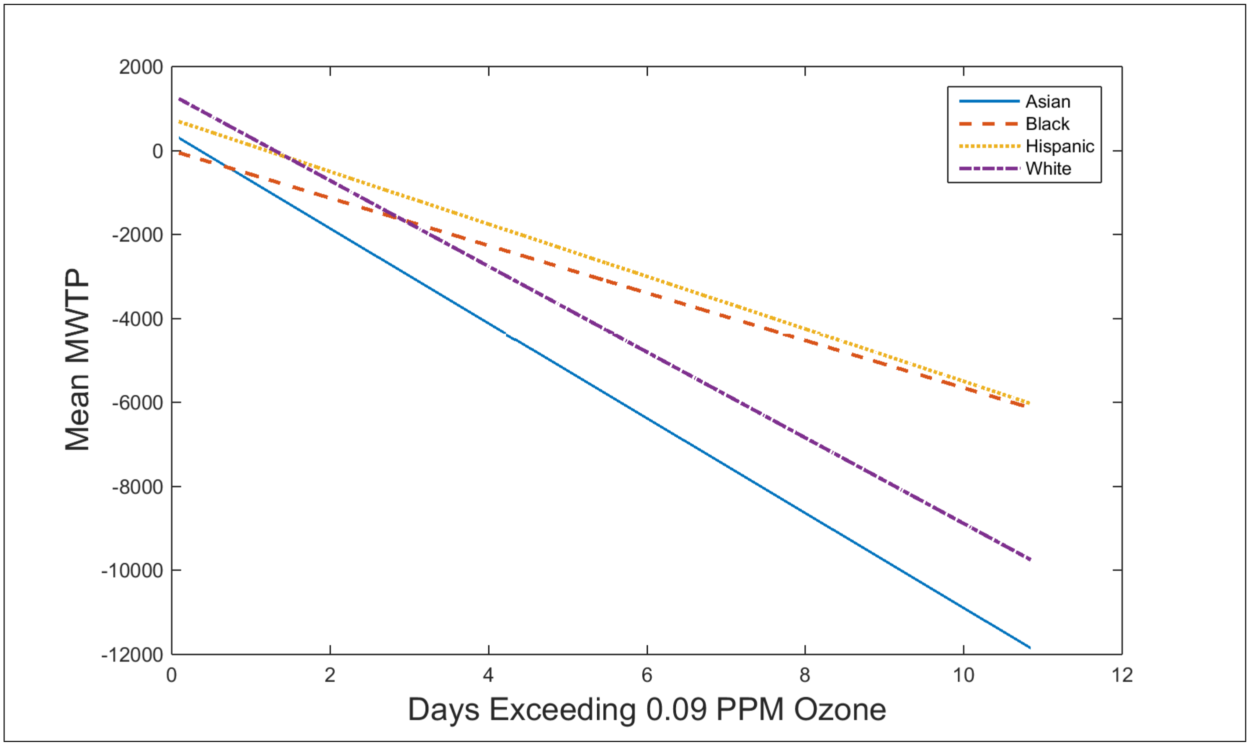

4.2.2 Demand Estimation with Three-Purchase Panel

In this case, we report the distribution of MWTP elasticities in Figure 8. With access to data describing three separate purchase occasions, we additionally report the distribution of elasticities with respect to income. This elasticity is calculated as $\hat{\alpha}_{5,i}\frac{I_{i,t}}{\hat{\rho}_{j^*(i),t}}$. As the quasi-linear utility function is linear in consumption, it is not ex-ante obvious whether this elasticity should be positive or negative. A negative elasticity would imply $\hat{\alpha}_{5,i} > 0$, meaning that ozone and consumption are complements and that for higher levels of consumption, the marginal (dis)utility of ozone is closer to zero. As can be seen in Figure 8, the income elasticities are centered around zero. The median value is -.0053.

4.2.3 Alternative Specifications - Inversion with Cross-Sectional Data

For the sake of comparison, we consider two alternative second-stage specifications that allow for the recovery of individual-specific preference parameters. These specifications

Figure 8: Distributions of MWTP Elasticities for Three-Purchase Panel

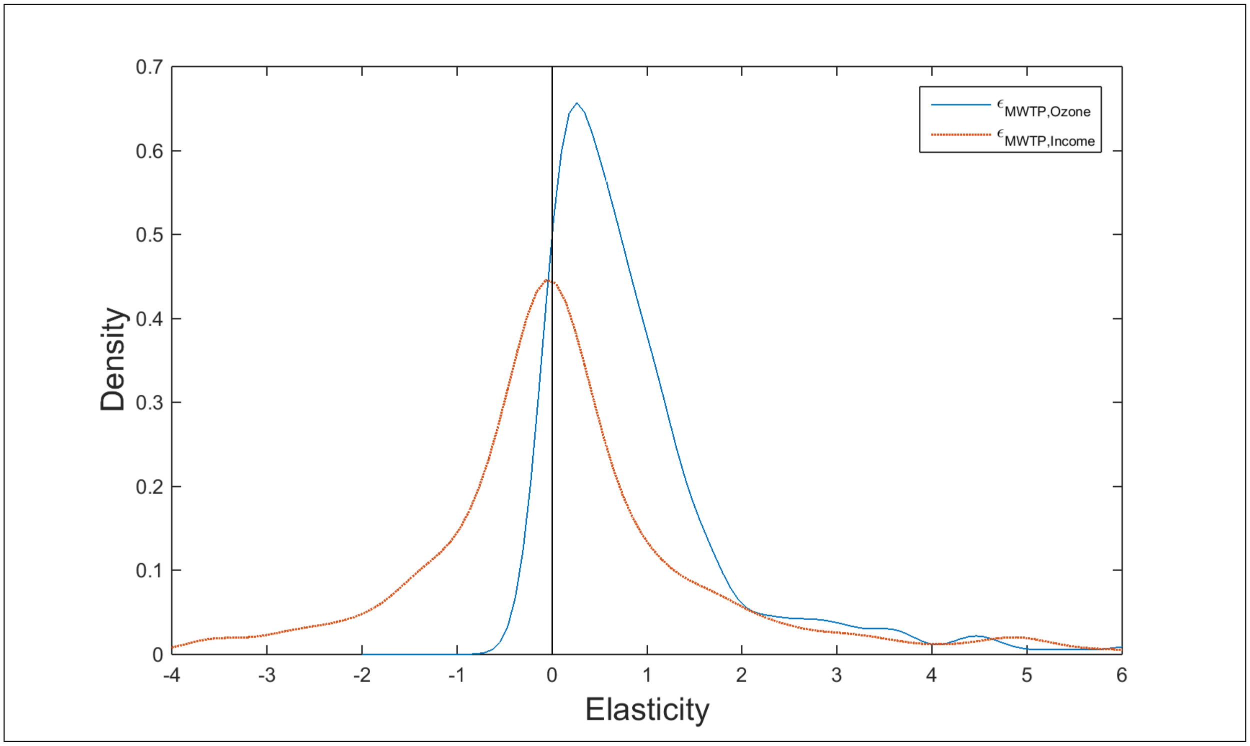

employ cross-sectional data and are commonly-used approaches to recovering the MWTP function. To implement these, we pool our data over time and treat our panel of both two- and three-purchase households as a cross-section.

The first alternative is to allow for heterogeneity in the intercept of the MWTP function, but specify that the MWTP function is flat.[31] As we know that the MWTP at the point of consumption is given by $\rho_{j^*(i),t}$, the entire MWTP function is simply recovered as $\hat{\rho}_{j^*(i),t}$. In other words, with flat MWTP functions the slope parameter $(\alpha_{4,i})$ is set to zero for all households. The distribution of estimated intercepts $(\hat{\alpha}_{3,i})$ is then given by Figure 5.

The second alternative is to use the functional form of the utility function to pin down the slope of the MWTP function in the spirit of Bajari and Benkard (2005). We

[31]As noted in Bishop and Timmins (2015), a flat MWTP function corresponds to either Bajari and Benkard (2005) with a linear utility function or the Rosen (1974) model with a linear utility function where the standard endogeneity problem is trivially solved by assuming that the MWTP function does not depend on Z_{jt}.

implement this by specifying an indirect utility function of:

$$V_{i,j,t} = \alpha_{0,i} + \alpha_{1,i}X_j + \frac{\alpha_{2,i}}{2}X_j^2 + \gamma_i exp(Z_{j,t}) + (I_{i,t} - R_{j,t}) \tag{19}$$

which yields a MWTP function of $\gamma_i exp(Z_{j,t})$.[32] An estimate of the parameter γ_i may then be recovered as:

$$\hat{\gamma}_i = \frac{\hat{\rho}_{j^*(i),t}}{exp(Z_{j^*(i),t})} \tag{20}$$

While this specification does allow MWTP to vary with both ozone and income, it does so in a highly restrictive way. In particular, the semi-elasticity of MWTP with respect to $Z_{j,t}$ (i.e., $\frac{\partial log(MWTP(Z_{j,t}))}{\partial Z_{j,t}}$) is restricted to be 1. Accordingly, we refer to this model as the unitary semi-elasticity model. That such a restriction is required is not surprising; a curve describing an individual's MWTP function is being identified from a single data point. In this unitary semi-elasticity model, the slope of the MWTP function (evaluated at the point of exposure) is, by construction, equal to the value of MWTP (evaluated at the point of exposure). As previously stated, these local values of MWTP may be seen in Figure 5. For comparison to our panel-data results, we could consider the elasticity of MWTP with respect to exposure level; in this case, however, they are essentially uninformative, as the elasticity is simply defined as $Z_{j,t}$.

5 Estimating the Welfare Effects of a Non-marginal Change in Air Quality

To illustrate the importance of heterogeneous (and downward-sloping) MWTP functions, we consider the valuation of a non-marginal change in ozone exposure. Correctly estimating heterogeneity in the slope of the MWTP function is critical from a policy perspective as it is important to know if those who are exposed to larger changes in ozone are those who are more or less sensitive to these changes. In other words, do the households exposed to the largest changes have the steeper or flatter MWTP functions? This question is particularly important given that households likely sort based on their preferences for ozone.

[32]Note that a typical specification for "goods" would be $V_{i,j,t} = \alpha_{0,i} + \alpha_{1,i}X_j + \alpha_{2,i}X_j^2 + \gamma_i log(Z_{j,t}) + (I_{i,t} - R_{j,t})$. If Z is a "good" (i.e., $\gamma > 0$), then the MWTP function will be downward sloping. However, if Z is a "bad" (i.e., $\gamma < 0$), then the MWTP function will be upward sloping, which itself highlights the restrictiveness of using functional form to recover the slope of the MWTP function.

In this analysis, we consider the welfare effects associated with increasing each household's exposure to ozone by 33% of their existing exposure level. Given this framework, those with the largest exposure will receive the largest change in ozone.[33] To put this into context, a 33% increase implies that the average change in ozone exposure is one-half of the median within-year standard deviation of ozone.[34]

We first present the welfare results from our panel-data-estimated MWTP functions. We then compare these with the results from the flat MWTP functions to illustrate the importance of allowing for downward-sloping MWTP curves. Finally, we compare our results with those from the unitary semi-elastic MWTP functions to highlight the importance of estimating, rather than specifying, heterogeneity in MWTP slopes.[35]

Using our model with heterogeneous slopes, we find that the mean willingness to pay to avoid a 33% increase in ozone is \$953.80. There is considerable heterogeneity as evidenced by the interquartile range of \$1,064.25. The full distribution of welfare effects is shown in Figure 9.[36]

It is well-accepted in the literature that using marginal estimates to value non-marginal changes leads to biased welfare estimates. This is effectively what is being done in the flat MWTP case.[37] In our context, this bias would lead to smaller estimated welfare effects. We find this: the welfare effects are considerable smaller under the assumption of flat MWTP curves. The mean welfare loss is only \$796.30.[38]

Less well-studied in the literature is the importance of allowing for heterogeneity in the slope of MWTP functions (and the importance of correctly estimating this heterogeneity). To illustrate the importance of estimating heterogeneity, we compare our results with

[33] One could consider an alternative experiment where those with the lowest current levels experienced the largest change. The concepts underlying the importance of heterogeneity would similarly apply in that context.

[34] The within-year standard deviations for the 13 years are: 2.0671, 2.0533, 1.6361, 2.3547, 2.9686, 2.5915, 2.6309, 2.0909, 2.3211, 1.1030, 1.6390, 1.7795, and 1.4282. The median of these numbers is 2.0671 and a 33% increase implies an average change of 1.0303.

[35] For all specifications, we use the full sample of households (those that purchase two and three times), dropping only those households that have slopes below the 1st percentile in the distribution of $\hat{\alpha}_{4,i}$. Results are robust to this choice.

[36] The median willingness to pay to avoid the 33% increase is \$699.78.

[37] This bias exists even when one allows for heterogeneity in the intercepts, as done here. Further restricting the MWTP to have a common intercept would add an additional bias as, even in the case of marginal changes, the level of the MWTP would be incorrectly estimated.

[38] The median willingness to pay to avoid the 33% increase is \$589.18 and the interquartile range of the distribution is \$830.01. The full distribution of welfare effects is shown in Figure A.2 in the Appendix.

Figure 9: Welfare Costs of a Non-Marginal (33%) Increase in Ozone

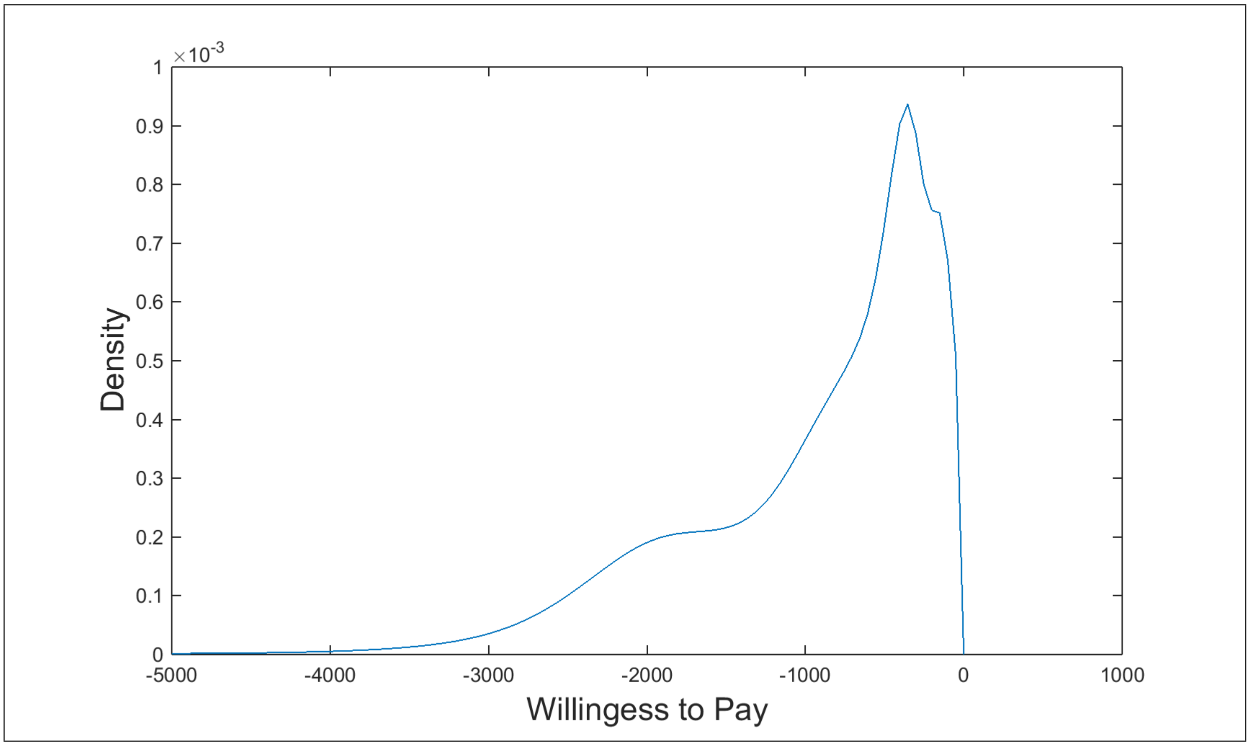

a Bajari-Benkard-style estimator where the MWTP function has a unitary semi-elasticity and is given by $\gamma_i exp(Z_{j,t})$.

In both our framework and the unitary semi-elasticity framework, the level of the MWTP function (at the point of consumption) is recovered by the first-order condition and the price gradient. The key distinction between the two approaches is in the recovery of the slope of the MWTP function going through that point.[39] In the unitary semi-elasticity framework, the slope of the MWTP function is dictated by the functional form of the utility function. In contrast, this paper shows how one can use panel data to estimate the slope of the MWTP function. Intuitively, households' price sensitivity is observed directly in the data: how much does the household change their ozone exposure when the implicit price of ozone changes?

Turning to the estimates, the unitary semi-elasticity model yields very different welfare estimates than our panel-data estimates. The mean welfare cost is more than

[39]It is also true for the flat-MWTP model that the level of the MWTP function (at the point of consumption) is recovered by the first-order condition and the price gradient.

twice as high at $2,437.22. The interquartile range is also much higher at $2,072.86.[40] To understand these differences, it is worth revisiting what dictates the slopes in the unitary semi-elasticity case. In the data, when the price gradient lies further from zero (i.e., households receive a bigger compensation for exposure to ozone), average exposure to ozone is larger. While this data pattern is to be expected, the framework of the unitary semi-elasticity model dictates that those who are exposed to the most ozone have steeper MWTP curves on average. This can be seen in the correlation between the slope of the MWTP function (evaluated at ozone exposure) and the amount of ozone exposure which is -0.13.[41] With a policy that proportionally changes ozone, the unitary semi-elasticity model assigns the most sensitive (steepest) MWTP curves to those who experience the greatest change.[42] This results in estimated welfare effects that are considerably higher on average and more dispersed.

In contrast, using panel data to estimate the slopes yields the (more sensible) result that those who consume the least amount of ozone are also typically those who are the most sensitive (i.e., those with the steepest MWTP slopes). With our estimated slopes, the correlation between the slope of the MWTP function and the amount of ozone exposure is +0.31.

Finally, we evaluate a 33% reduction in ozone with results following a similar, yet opposite, pattern: the welfare effects associated with this change are now positive and considerably smaller for the unitary semi-elasticity case. In our panel-data framework with estimated slopes, the mean welfare benefit is $638.79. In the unitary semi-elasticity framework with assigned slopes, the mean welfare benefit is only $429.76. By construction, the mean welfare benefit in the flat MWTP framework is simply the negative of the cost for the 33%-increase in ozone, i.e., $796.30.

[40] The median willingness to pay to avoid the 33% increase is $944.69. The full distribution of welfare effects is shown in Figure A.3 in the Appendix.

[41] As the price gradients are upward sloping, the within-year correlation between MWTP slope and ozone exposure is positive. However, in the overall data, this within-year positive correlation is outweighed by the strong between-year negative correlation between MWTP slope and ozone.

[42] If one were to work with a "good", instead of a "bad", one could adopt a unitary elasticity model like Cobb-Douglas, instead of the unitary semi-elasticity model estimated here. In that case, the core principal that MWTP slopes are identified solely by functional form would still hold. However, whether those consuming the most had steeper or flatter MWTP functions would depend on the estimated price gradients.

6 Conclusion

We show in this paper how access to panel data allows the econometrician to observe each household's sensitivity to changes in implicit prices, thus recovering a fully-heterogeneous estimate of the MWTP function. Applying this method to valuing a non-marginal change in ozone exposure in the Bay Area of California, we find evidence of large welfare effects. Importantly, we find serious implications for this valuation exercise when we apply existing methods in the literature: (i) assuming flat MWTP functions and (ii) assigning MWTP slopes based on functional-form assumptions. As one would expect, the assumption of flat MWTP functions biases the welfare costs associated with a proportional increase in ozone toward zero. Potentially less-expected is that the method of assuming MWTP-function slopes (via functional-form) determines the empirical correlation between ozone-exposure level and assigned slope. In our application, this is seen in the steepest slopes being assigned to those households with the highest levels of consumption, thus biasing the welfare costs associated with a proportional increase in ozone away from zero.

Importantly, the panel-data-based estimation approach described in this paper is increasingly implementable with the growing availability of panel datasets that match buyer attributes to the prices and characteristics of the properties they purchase. In addition, it is applicable to a wide variety of non-marginal policy analyses, including the welfare implications of changing the levels of other local pollutants or changing the levels of local school quality or crime rates. Importantly, this method allows for the recovery of fully-heterogeneous MWTP functions and for the analysis of non-marginal policy changes without encountering the well-known endogeneity issues with the recovery of hedonic demand functions.

7 References

Bajari, P. and L. Benkard (2005). "Demand Estimation With Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach." *Journal of Political Economy.* 113(6):1239-1276.

Banzhaf, S., (2015). "Panel-Data Hedonics: Rosen's First Stage and Difference-in-Differences as 'Sufficient Statistics."' *mimeo.*

Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2015). "A Dynamic Model of Demand for Houses and Neighborhoods." forthcoming, *Econometrica.*

Bartik, T.J. (1987). "The Estimation of Demand Parameters in Hedonic Price Models." *Journal of Political Economy.* 95(1):81-88.

Bishop, K. and Murphy, A. (2015). "Valuing Time-Varying Attributes using the Hedonic Model: When is a Dynamic Approach Necessary?" *Mimeo.*

Bishop, K. and Timmins, C., (2015). "Estimating the Marginal Willingness to Pay Function Without Instrumental Variables" *Mimeo.*

Brown and Rosen (1982). "On the Estimation of Structural Hedonic Price Models." *Econometrica.* 50(3):765-768.

Case, K.E., and Schiller, R.J. (1989). "The Efficiency of the Market for Single-Family Homes." *American Economic Review.* 79(1):125-137.

Ekeland, I., Heckman, J., and Nesheim, L. (2004). "Identification and Estimation of Hedonic Models." *Journal of Political Economy.* 112(1) S60–S109

Epple, D. (1987). "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products." *Journal of Political Economy.* 95(1):59-80.

Fan, J. and Gijbels, I. (1996). "Local Polynomial Modeling and Its Applications." In *Monographs on Statistics and Applied Probability* (66). Chapman and Hall. London, U.K.

Lee, Y. and D. Mukherjee (2014). "Nonparametric Estimation of the Marginal Effect in Fixed-Effect Panel Data Models: An Application on the Environmental Kuznets Curve." *Mimeo*

Mendelsohn, R. (1985). "Identifying Structural Equations with Single Market Data." *Review of Economics and Statistics*. 67(3):525-529.

Murray, M. (1975). "The Distribution of Tenant Benefits in Public Housing." *Econometrica*. 43:771-788.

________ (1983). "Mythical Demands and Mythical Supplies for Proper Estimation of Rosen's Hedonic Price Model." *Journal of Urban Economics*. 14:327-337.

Poterba, J. (1984) "Tax Subsidies to Owner-Occupied Housing: An Asset Market Approach." *Quarterly Journal of Economics*. 99(4) 729–752

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

Appendix

A.1 Price Gradients

Figure A.1: Price Gradients by Year

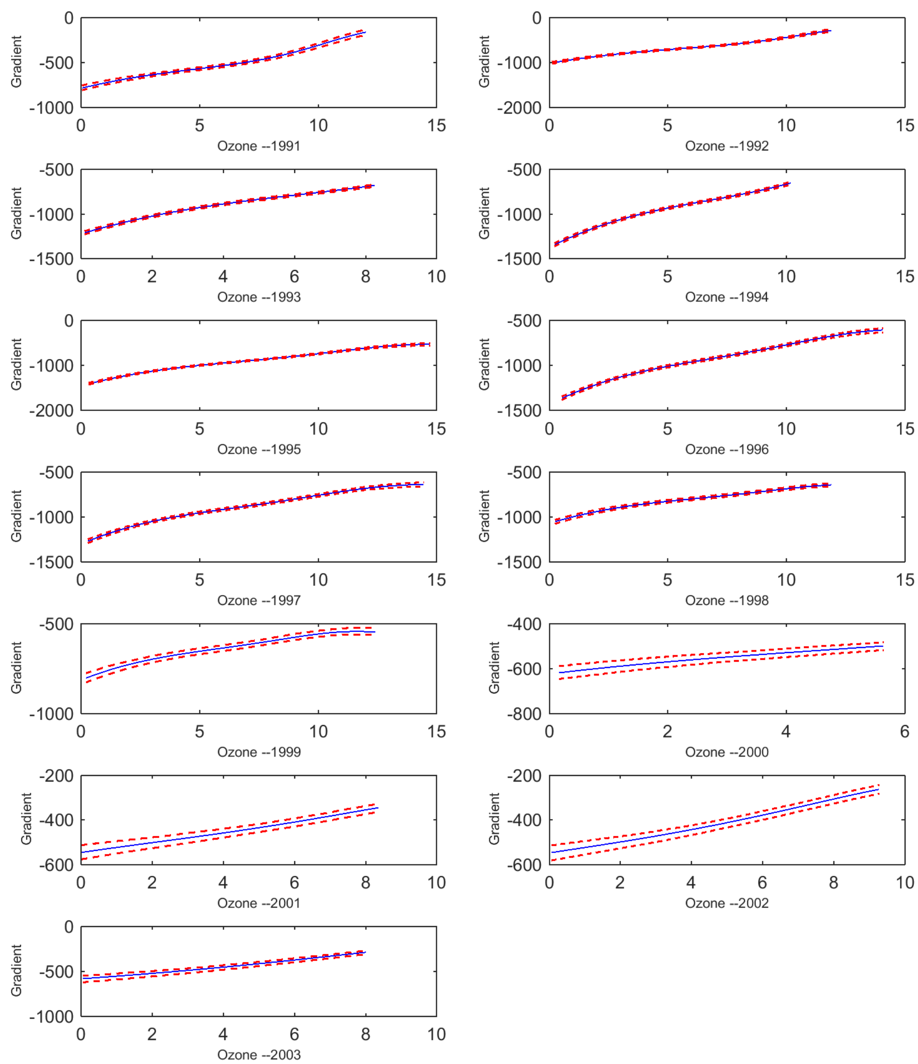

The 95% confidence intervals were generated using a bootstrap procedure with 250 draws.

A.2 Demand Estimation with Three-Purchase Panel

The first order condition, (17), will hold in each of the three periods the household is observed making a housing choice decision, which yields the following three equations:

$$\begin{aligned} \alpha_{3,i} + \alpha_{4,i} Z_{j^*(i),1} + \alpha_{5,i}(I_{i,1} - R_{j^*(i),1}) - \rho_{j^*(i),1}(1 + \alpha_{5,i} Z_{j^*(i),1}) &= 0 \\ \alpha_{3,i} + \alpha_{4,i} Z_{j^*(i),2} + \alpha_{5,i}(I_{i,2} - R_{j^*(i),2}) - \rho_{j^*(i),2}(1 + \alpha_{5,i} Z_{j^*(i),2}) &= 0 \\ \alpha_{3,i} + \alpha_{4,i} Z_{j^*(i),3} + \alpha_{5,i}(I_{i,3} - R_{j^*(i),3}) - \rho_{j^*(i),3}(1 + \alpha_{5,i} Z_{j^*(i),3}) &= 0 \end{aligned} \tag{A.1}$$

Solving this system provides us with the following equations for recovering household i's preference parameters:

$$\begin{aligned} \hat{\alpha}_{4,i} &= \frac{\theta_{j^*(i),3}(\hat{\rho}_{j^*(i),1} - \hat{\rho}_{j^*(i),2}) - \theta_{j^*(i),2}(\hat{\rho}_{j^*(i),1} - \hat{\rho}_{j^*(i),3})}{\theta_{j^*(i),3}(Z_{j^*(i),1} - Z_{j^*(i),2}) - \theta_{j^*(i),2}(Z_{j^*(i),1} - Z_{j^*(i),3})} \\ \hat{\alpha}_{5,i} &= \frac{\hat{\alpha}_{4,i}(Z_{j^*(i),1} - Z_{j^*(i),3}) - \hat{\rho}_{j^*(i),1} + \hat{\rho}_{j^*(i),3}}{\hat{\rho}_{j^*(i),1} Z_{j^*(i),1} - \hat{\rho}_{j^*(i),3} Z_{j^*(i),3} + (I_{i,3} - R_{j^*(i),3}) - (I_{i,1} - R_{j^*(i),1})} \\ \hat{\alpha}_{3,i} &= \hat{\rho}_{j^*(i),1}(1 + \alpha_{5,i} Z_{j^*(i),1}) - \hat{\alpha}_{4,i} Z_{j^*(i),1} - \hat{\alpha}_{5,i}(I_{i,1} - R_{j^*(i),1}) \end{aligned} \tag{A.2}$$

where:

$$\begin{aligned} \theta_{j^*(i),2} &= \hat{\rho}_{j^*(i),1} Z_{j^*(i),1} - \hat{\rho}_{j^*(i),2} Z_{j^*(i),2} + (I_{i,2} - R_{j^*(i),2}) - (I_{i,1} - R_{j^*(i),1}) \\ \theta_{j^*(i),3} &= \hat{\rho}_{j^*(i),1} Z_{j^*(i),1} - \hat{\rho}_{j^*(i),3} Z_{j^*(i),3} + (I_{i,3} - R_{j^*(i),3}) - (I_{i,1} - R_{j^*(i),1}) \end{aligned} \tag{A.3}$$

A.3 Change in Ozone Welfare Effects

Figure A.2: Welfare Costs of a Non-Marginal Increase Under Flat MWTP

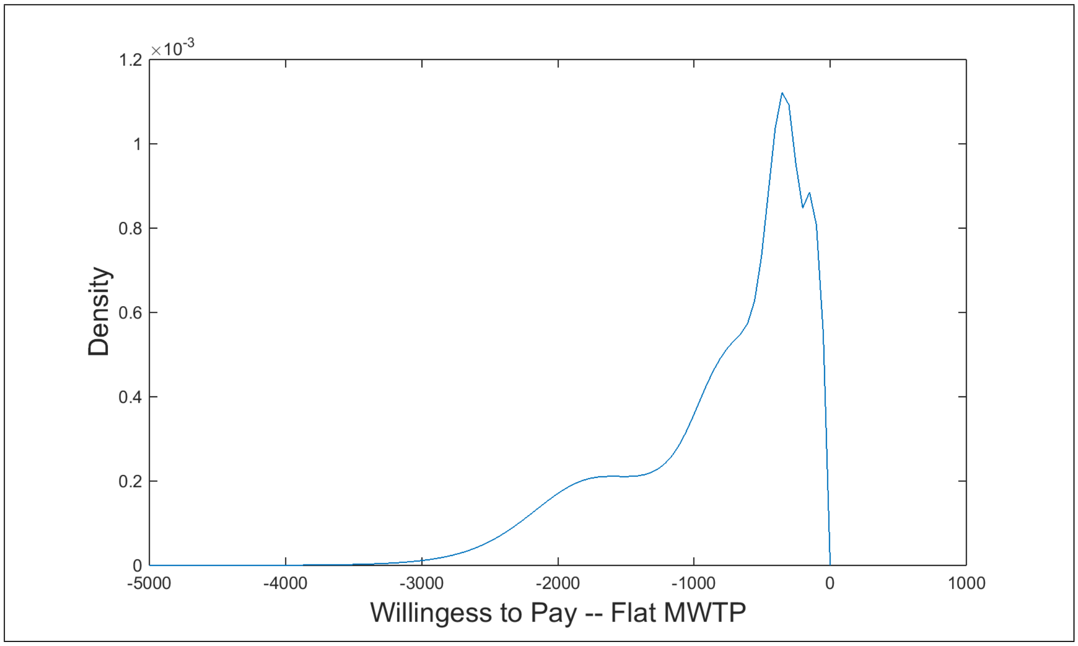

Figure A.3: Welfare Costs of a Non-Marginal Increase Under Unitary Semi-Elastic MWTP