Using Panel Data to Easily Estimate Hedonic Demand Functions

Kelly C. Bishop, Christopher Timmins

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 preference heterogeneity for a given amenity. In particular, in the face of a changing price schedule, observing each individual twice is sufficient to recover a linear demand function separately for each individual, with no additional restrictions. Constructing a rich panel data set of buyers, 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 a local estimate of marginal willingness to pay, is important. Second, we find evidence of considerable heterogeneity, which is important from a policy perspective; our data-driven estimates of the welfare effects associated with a nonmarginal change in air quality differ substantially from those recovered using the existing approaches to welfare estimation.

JEL Codes: Q50, Q51, R21, R23

Keywords: Hedonic demand, Ozone, Valuation

MANY APPLICATIONS OF THE HEDONIC MODEL seek to value nonmarginal 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 that 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.

Kelly C. Bishop (corresponding author) is at Arizona State University, Tempe, AZ 85287, USA (kelly.bishop@asu.edu). Christopher Timmins is at Duke University, Box 90097, Durham, NC 27708, USA (christopher.timmins@duke.edu). 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, St. Louis University, Econometric Society summer meetings, University of Calgary, and University of Toronto for their helpful comments and suggestions.

Received February 14, 2016; Accepted December 27, 2017; Published online April 10, 2018.

JAERE, volume 5, number 3. 2333-5955/2018/0503-0001$10.00 http://dx.doi.org/10.1086/696981

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 that when the hedonic price function is nonlinear 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 nonlinear budget constraint creates a difficult endogeneity problem when using statistical inference to recover the parameters describing those preferences.[1] Typically, instrumental-variable approaches in this literature have relied upon questionable exclusion restrictions; for example, certain sociodemographic variables enter directly into the MWTP function while others do not and the excluded variables can be used as instruments for endogenous attribute levels. 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 a 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 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, that is, 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.

Our method for recovering the MWTP function is quite intuitive: observing the exact same household make purchase decisions in multiple geographic markets or

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 et al. (2004) clearly demonstrate the conditions under which the MWTP function is identified. Bishop and Timmins (2018) build upon the insights of Ekeland et al. 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.

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 data set. Figure 1 describes this approach for a given 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 function for each household in the data. Naturally, the MWTP function that we recover with exactly two observed purchases is linear. 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. Additionally, we show how these preference coefficients may be decomposed to isolate the effect of fixed demographic characteristics, such as race. While we recover a linear MWTP with two observations per household, we show that with three or more observations for each household, we can estimate the effect of time-varying demographic characteristics, such as income. With three or more observations per household, we could alternatively allow for greater flexibility in the

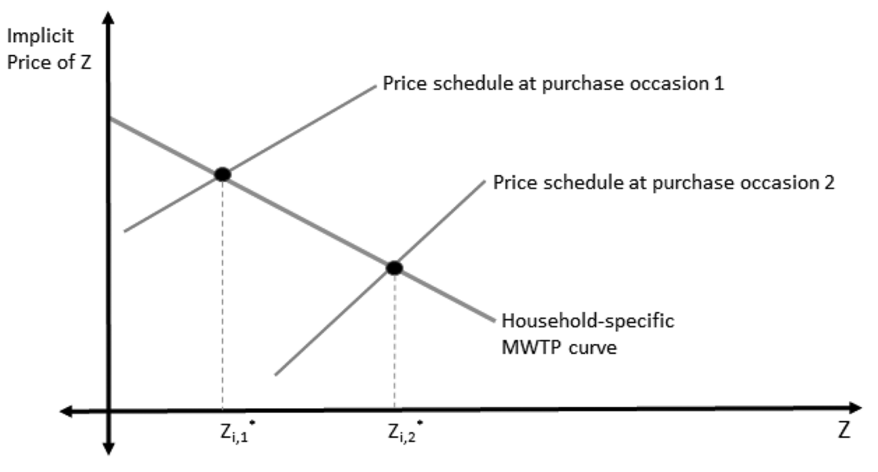

Figure 1. Recovering hedonic demand functions using panel data

shape of the MWTP curve, such as allowing the function to be quadratic in Z or allowing nonseparability between Z and other housing characteristics.

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 data set describing real estate transactions and the attributes of the associated buyers over the 13-year period 1991–2003. To create this data set, we first isolate and match individual households over time and then merge the demographic characteristics provided on mortgage applications from 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 number of days exceeding the state of California's 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 nonparametric regression to recover a flexible set of time-varying hedonic price gradients which includes house-level fixed effects. 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 value of this elasticity is 0.69, implying that it is important to recover the entire function when considering nonmarginal changes. We also find considerable heterogeneity across households in the MWTP to pay to avoid ozone pollution with an interquartile range of 0.89 in this 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 who are exposed to larger policy-induced changes in amenities may be those who are 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 annual welfare costs associated with a nonmarginal increase in ozone. We find that the mean willingness to pay to avoid a 33% increase in ozone is \$1,021, the median is \$770, and the interquartile range is \$1,106. We 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 (i.e., perfectly elastic) MWTP functions. Using these flat MWTP functions returns the expected result of smaller mean annual welfare costs: the mean cost is \$848 and the median cost is \$636. 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, the estimated annual welfare costs are very close to the model with flat MWTP functions: the mean is $909 and the median is $652. Furthermore, the functional-form based approach used here recovers a counterintuitive negative correlation between ozone exposure and the slope of the MWTP function; the steepest MWTP curves are assigned to those households that 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 ozone exposure and the slope of the MWTP function.

This paper proceeds as follows. Section 1 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 data set and its summary statistics are discussed in section 2. Section 3 presents our results, and section 4 applies our MWTP-function results to measure the willingness to pay to avoid a nonmarginal change in ozone levels. Finally, section 5 concludes.

1. MODEL

In this section, we describe our panel-data-driven approach to recovering the structural parameters of the hedonic model. First, we discuss the nonparametric, 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 that 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.

1.1. A Nonparametric Fixed-Effects Approach to Recovering the Hedonic Gradient

In the first stage of the model, we estimate the hedonic price function, $P_{j,t}$ which relates the price of house j transacted in period t to its attributes: both those that vary over time, $Z_{j,t}$, and those that do not, X_j. The elements of X_j may be observed or unobserved. 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.[2] Additionally, concern must be paid to the potential bias caused by omitted variables. 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.

2. See Ekeland et al. (2004) for a discussion of the shape of equilibrium hedonic price functions.

Following Fan and Gijbels (1996) and Lee and Mukherjee (2014), we adopt a nonparametric first-difference approach for the estimation of the gradient and begin by writing down a flexible representation of the hedonic price function:[3]

$$P_{j,t} = f\left(Z_{j,t}\right) + X_j'\theta + \nu_{j,t}, \tag{1}$$

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

In a series of Monte Carlo experiments, Kuminoff et al. (2010) show that including spatial fixed effects is the preferred way to deal with (potentially correlated) unobservables at the neighborhood level. Here, as we have repeat-sales data, we are controlling for these unobservables at the level of the house.

We take a first-order Taylor series expansion of $f(\cdot)$ around a vector z, which has the same dimension as $Z_{j,t}$.[4]

$$P_{j,t} = f(z) + \left(Z_{j,t} - z\right)'f'(z) + X_j'\theta + \nu_{j,t}. \tag{2}$$

We denote each property's year of prior sale by t_{-1} and rewrite equation (2) for this prior sale:[5]

$$P_{j,t_{-1}} = f(z) + \left(Z_{j,t_{-1}} - z\right)'f'(z) + X_j'\theta + \nu_{j,t_{-1}}. \tag{3}$$

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, $f(z)$.[6]

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

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

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

3. Bajari and Kahn (2005) apply a similar framework, but without the first differencing. See also Clapp (2003), Parmeter et al. (2007), and Heckman et al. (2010) for discussions of nonparametric hedonic price function estimation.

4. The remainder term associated with the Taylor expansion is ignored.

5. This is expanded around the same vector z.

6. Losing $f(\cdot)$ from this expression does not pose a problem, as our interest is only in recovering the hedonic gradient, that is, the slope of the hedonic price function, which is represented nonparametrically by $f'(\cdot)$.

Lee and Mukherjee (2014) show that $\beta(z)$ (i.e., the slope of the hedonic price function at $Z_{j,t} = z$) may be recovered using the following least-squares minimization procedure:

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

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, z. For our estimation, $K_h(\cdot)$ is given by the Gaussian kernel:

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

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

In practice, this allows us to recover an estimate of $\beta(z)$ for all observed values of $Z_{j,t}$. We therefore end up with a (potentially) different estimate of $\beta(z)$ 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 $Z_{j,t}$.

We now discuss which features of the data identify $\beta(z)$ in equation (5). As we are differencing out the house-specific, time-invariant characteristics, $\beta(z)$ is identified by the remaining, within-house variation in housing price, year of sale, and ozone. Given the nonparametric model, price is locally linear in both year and ozone, that is, in the vector $Z_{j,t}$. This means that for each value of z, $\beta(z)$ is identified using within-house variation from houses with values of $Z_{j,t}$ and $Z_{j,t_{-1}}$ "close" to z.

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

$$\beta(z) = \left(\tilde{Z}' W_z \tilde{Z}\right)^{-1} \tilde{Z}' W_z \tilde{P}, \tag{8}$$

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

Finally, we note that as t_{-1} is not the previous year, but rather the year of the house's prior sale, the time gap between house sales used to difference prices in (4) varies across houses. As such, the data are not a standard panel as the intervals between observations are random.[7]

7. Other recent work that estimates nonparametric models with panel data includes Henderson et al. (2008), Qian and Wang (2012), and Su and Lu (2013).

1.2. 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 nonhousing consumption ($C_{i,t}$), choosing a house j with attributes (X_j, $Z_{j,t}$), in period t:

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

where we have normalized the coefficient on consumption to one. Preference parameters (i.e., the α_i) are assumed to be stable over time. We consider alternative utility functions in the appendix, available online.

Household i faces a budget constraint, $C_{i,t} + R_{j,t} \leq I_{i,t}$, where $I_{i,t}$ denotes income and $R_{j,t}$ denotes the imputed annual rent or housing expenditure associated with house j. In practice, we calculate this figure as 5% of the observed transaction price.[8]

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

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

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:

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

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

$$MWTP\left(Z_{j,t}\right) = \alpha_{3,i} + \alpha_{4,i}Z_{j,t}. \quad (12)$$

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.[9] Fortunately, we observe household

---

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

9. 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 $(\partial R_{j,t})/(\partial Z_{j,t})$ 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.

i in panel data on (at least) two occasions. For households observed twice, we then have two 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 nonparametric 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, that is, the household-specific solutions for both the intercept and the slope of the MWTP function.[10] Respectively, these are given by:

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

As the second stage is an inversion, rather than an estimation, the only sampling variance of the estimates comes from the first stage.[11] While the second stage does not add sampling variation, the standard errors will be a function of the difference between $Z_{j,1}$ and $Z_{j,2}$ observed in the second-stage data. As can be seen in the formulas for the variance of $\hat{\alpha}_{3,i}$ and $\hat{\alpha}_{4,i}$ in equation (15), households that consume similar amounts of Z in both periods will have larger standard errors, all else equal. In practice, we bootstrap all standard errors using 1,000 draws.

$$\begin{aligned}
\operatorname{var}(\hat{\alpha}_{3,i}) &= \frac{\operatorname{var}\left(\hat{\rho}_{j^*(i),2}\right) Z^2_{j^*(i),1} + \operatorname{var}\left(\hat{\rho}_{j^*(i),1}\right) Z^2_{j^*(i),2} - 2\operatorname{covar}\left(\hat{\rho}_{j^*(i),2}, \hat{\rho}_{j^*(i),1}\right) Z_{j^*(i),1} Z_{j^*(i),2}}{\left(Z_{j^*(i),1} - Z_{j^*(i),2}\right)^2} \\
\operatorname{var}(\hat{\alpha}_{4,i}) &= \frac{\operatorname{var}\left(\hat{\rho}_{j^*(i),1}\right) + \operatorname{var}\left(\hat{\rho}_{j^*(i),2}\right) - 2\operatorname{covar}\left(\hat{\rho}_{j^*(i),1}, \hat{\rho}_{j^*(i),2}\right)}{\left(Z_{j^*(i),1} - Z_{j^*(i),2}\right)^2}.
\end{aligned} \tag{15}$$

---

10. Having two equations is a result of observing the household purchase twice. Having two unknowns reflects the assumption that the MWTP function is linear. As such, identifying (completely heterogeneous) nonlinear MWTP functions would require observing the household make more than two purchases. See Harberger (1971) and Banzhaf (2015) for a discussion about how using a linear MWTP function provides a second-order approximation to nonmarginal welfare measures for any possible constant demand function.

11. The price function is very precisely estimated (see fig. A1). Consequently, standard errors for the second-stage MWTP estimates and welfare calculations are very small, as can be seen in figures A2–A6.

It is possible to determine how households' MWTP functions differ systematically with fixed household 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:

$$\hat{\alpha}_{3,i} = \delta_{3,0} + A_i'\delta_{3,1} + \varsigma_{3,i} \qquad \hat{\alpha}_{4,i} = \delta_{4,0} + A_i'\delta_{4,1} + \varsigma_{4,i}. \tag{16}$$

While estimating a linear MWTP is an improvement over the common alternative of estimating a flat (i.e., perfectly elastic) MWTP function, it remains the case that with two observations per household, we are restricted to recovering a linear (but fully heterogeneous) MWTP function. Having access to more than two observations per household would allow the assumption of linearity to be weakened. For example, with three observations per household, one could recover a quadratic MWTP function.

With three (or more) observations per household, one could alternatively allow for richer forms of utility. One example would be to relax the assumption that preferences do not change over time and allow for time-varying MWTP functions by estimating a coefficient on a time-varying attribute, such as income. Another example would be to allow for nonseparability between the amenity of interest, Z_{jt}, and another (time-invariant) housing characteristic in X_j. We elaborate on both of these examples (and present results from the estimation of these alternative specifications) in the online appendix.

Finally, we note that this approach relies critically on having access to data where households can be observed making more than one purchase. As such, the simplicity of connecting the dots comes with higher data requirements.

2. DATA

To implement our first-stage nonparametric regressions, we use data describing single-family housing transactions over the period 1991–2003 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 data set by combining information from the real estate transactions data set and a data set describing mortgage applicants' demographic characteristics obtained through the Home Mortgage Disclosure Act (HMDA).

2.1. 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

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 that could be solved for the four unknowns, $\gamma_{3,0,i}$, $\gamma_{3,1,i}$, $\gamma_{4,0,i}$, $\gamma_{4,1,i}$. 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.

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 of bedrooms.

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 rebuilds, 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 rebuild, 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).[13] 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 1% 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 missing.[14]

Finally, we restrict our analysis to properties with multiple sales over the 13-year period.[15] This yields a final sample of 277,011 transactions (i.e., property-year observations) comprising 126,227 unique properties. Table 1 describes the data.

2.2. Buyer Characteristics Data

In order to implement the second stage of our estimator, we need to create a panel of households and their chosen properties over time.[16] This involves first identifying and matching households over time in the property transactions data set and, second, merging-in the attribute data from the HMDA data set. This is possible as we have

13. 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.

14. Although we use a repeat-sales estimation approach, we make some cuts on property characteristics to ensure that all single-family homes are trading in the same market: dropping properties where year built is pre-1850, lot size 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 10, number of total rooms is greater than 15, or number of stories is greater than three.

15. However, we drop properties that sell more than once within a calendar year or more than five times over the 13-year period.

16. While observing the characteristics of the buyers is not necessary for recovering the MWTP functions, observing them allows us to decompose preferences by race and allows us to check whether observable covariates are balanced between the one-purchase and multiple-purchase samples.

Table 1. Housing Transactions Data

	Full Sample n = 630,384		Repeat-Sales Sample n = 277,011	
Variable	Mean	SD	Mean	SD
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
No. bedrooms	3.20	.90	3.12	.88
No. bathrooms	2.12	.73	2.09	.69

buyer and seller names in the transaction record as well as common variables in both data sets.

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 data set 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]

The individual attribute data come from a data set on mortgage applications published through the Home Mortgage Disclosure Act of 1975. These data are publicly available and provide information on all mortgage applications filed in the Bay Area over the period of our sample.[20] Included are all applicants' race and gender, income,

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 suggests 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 suggests that only using within-county moves is appropriate.

19. If a household only purchases once 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 13-year sample period), we assume that he or she is a professional real-estate "flipper" and drop them from the sample. Similarly, we drop buyers who are observed buying more than once within a calendar year.

20. The HMDA data are available at https://www.ffiec.gov/hmda/hmdaproducts.htm.

Table 2. One-, Two-, and Three-Purchase Samples of Households

	One Purchase $n = 315{,}215$		Two Purchase $n = 6{,}057$		Three Purchase $n = 218$	
Variable	Mean	Median	Mean	Median	Mean	Median
Asian	.24	0	.26	0	.29	0
Black	.03	0	.03	0	.01	0
Hispanic	.11	0	.13	0	.15	0
White	.61	1	.58	1	.54	1
Income (in yr 2000 $)	114,464	96,014	125,779	103,476	132,934	109,224
Price (in yr 2000 $)	390,801	347,580	413,834	359,373	394,349	330,067

loan amount, lender name, and census tract of the property. We are able to merge the individual attributes in the HMDA data set to the buyers in the property transactions data set using the common variables of lender name, loan amount, transaction date, and census tract of the property.[21] We successfully match approximately two-thirds of households in the transactions sample to the raw HMDA sample.[22]

Table 2 provides statistics describing the one-purchase, two-purchase, and three-purchase buyers. It is important to note that our estimation only uses those who purchase two or three times over the sample period, which in principle, could result in a nonrandom sample if those who purchase more than once have different preferences. However, based on observable characteristics, the two- and three-purchase sample is very similar to the one-purchase sample. There are very minor differences in race as well as the price of housing. Mean income is approximately 10% higher in the two- and three-purchase sample, with a slightly smaller differential in median income.

2.3. Ozone Data

The ozone data we employ are taken from the California Air Resources Board.[23] We use yearly ozone data from 35 monitors in the nine counties of Alameda, Contra Costa,

21. See Bayer et al. (2016) for more details on this merge. Combining transactions data with HMDA has become more common in the literature (e.g., Diamond and McQuade 2016).

22. 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 is 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.

23. Publicly available at www.arb.ca.gov/adam/.

Marin, Napa, San Francisco, San Mateo, Santa Clara, Solano, and Sonoma over the period 1990–2004. In particular, we use the monitor data to construct property-specific measures of the number of days exceeding the 1-hour California standard (i.e., 90 parts per billion).

In addition to the ozone readings, the data set provides information on the "year coverage," or the percentage of time (during the relevant high-ozone season) each particular monitor was available and the geographic coordinates of each monitor.[24] Using this coverage variable, we drop monitors with less than 60% coverage in a given year (amounting to less than 4% of the available monitor-year observations).

Using the latitudinal and longitudinal coordinates of both the monitors and the properties, we use the "Great Circle" algorithm to compute the distance to all monitors from each property. We then create a 3-year weighted average for each property of all monitors' readings, weighting distance by one over distance squared.[25] 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 attainment and mandatory ozone reduction programs

24. Some monitors were opened or permanently closed during this time period.

25. We use a 3-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. A similar weighted approach was used in Chay and Greenstone (2005). See Bishop and Murphy (2018) for a discussion of using time-averaged amenities in hedonic analyses.

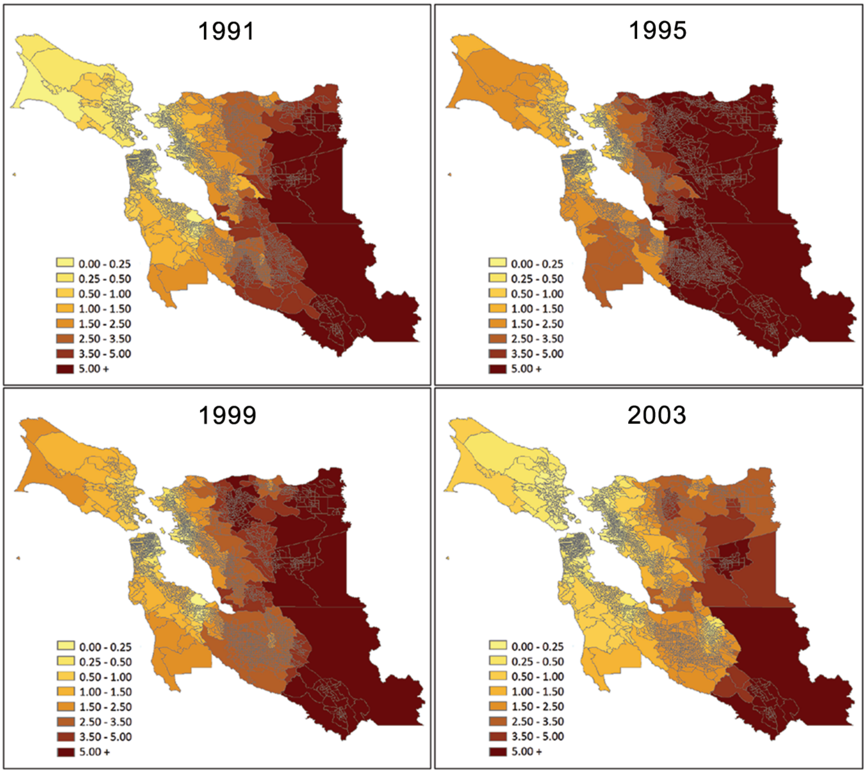

Figure 2. Ozone pollution in the Bay Area

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

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 (i.e., ozone exceedances) as our measure of ground-level ozone, as it is a good measure of the frequency of extreme pollution events that house buyers may be aware of. As 1-hour ozone concentration is a component of the Air Quality Index that is published online and in newspapers, another po-

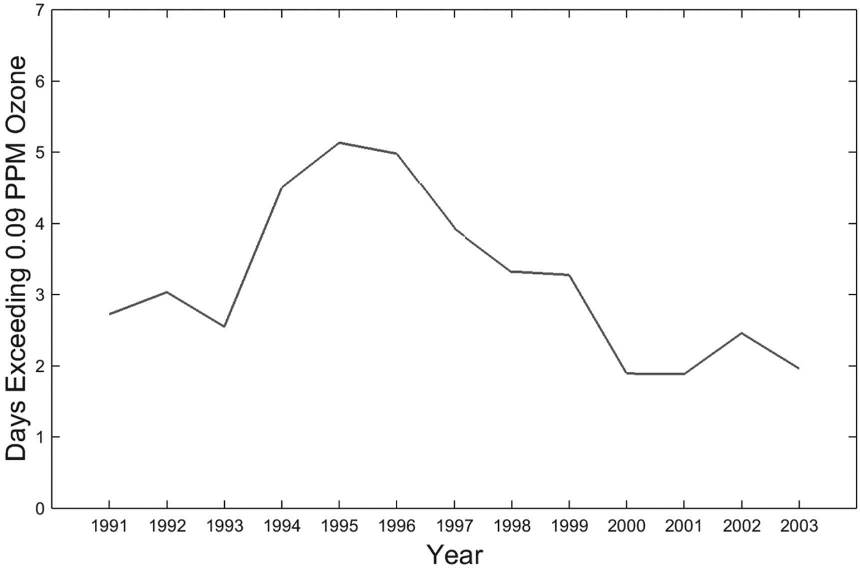

Figure 3. Average number of days exceeding the CA 1-hour standard

tential measure of ozone would be the yearly maximum 1-hour ozone concentration (which is also available through the California Air Resources Board). Therefore, in addition to estimating the model using the exceedance data, we also estimate the model using the maximum 1-hour concentration data. Estimation results are similar across the two measures of ozone and are presented in the appendix.

3. RESULTS

In this section, we describe our results. First, we illustrate the results of our nonparametric 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 6,059 households that are observed purchasing two houses in our data set to recover estimates of fully heterogeneous MWTP functions as described in section 1.2. We also use the sample of 218 households that buy three times to recover estimates of MWTP functions that are allowed to vary with nonhousing 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.

3.1. Results from the First-Stage Hedonic Regressions

The local linear estimation allows the estimate of the slope coefficient, $\beta(z)$, 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, as having monotonic gradients allows us to compare our results with approaches where the MWTP function is flat. Ideally, we would use cross-validation to choose these smoothing parameters, but the model in which the MWTP function is flat requires monotonicity to avoid violations of the second-order conditions for utility maximization. Thus, we oversmooth to achieve monotonicity.[26] An alternative approach to ensure monotonicity of the gradients would be to directly impose shape restrictions on the price function, that is, impose that the price function is monotonic, convex, and smooth.[27] However, the existing shape-restriction-based approaches are not applicable as we face a multivariate problem with a large number of observations.[28] In practice, we isolate the set of bandwidths over ozone and year that ensure monotonicity and choose the pair with the smallest bandwidth for ozone.[29] We present results under different bandwidth choices in the appendix and show that our results are robust to the choice of bandwidths.

We estimate hedonic gradients for each of the 13 years in our sample. For ease of exposition, we show the hedonic gradient for three of these years (1992, 1997, and 2002) in figure 4 and show all of the gradients with 99% confidence intervals in figure A1 (figs. A1–A10 are available online in the appendix).[30] 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 upward sloping (i.e., the compensating reduction in price gets smaller/closer to zero at higher pollution levels). This suggests that households may sort based on preferences

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.

27. See Henderson and Parmeter (2015) for a comprehensive discussion of shape-based restrictions in nonparametric models.

28. Using constraint-weighted bootstrapping (e.g., Hall and Huang 2001; Du et al. 2013) to ensure monotonicity of the price gradient would require solving a quadratic program for each of the approximately 160,000 observations in the data. See Ryan and Wales (2000) for an illustration of a case where enforcing a constraint to hold at a single, well-chosen point can result in nearly all of the observations satisfying that constraint.

29. This is achieved by setting the bandwidth over ozone to 4.01 and the bandwidth over year to 3.35. We drop outlier households that choose quantities of ozone in 1999 that exceed the 99th percentile of ozone in that year, as the far right tail of the 1999 gradient experienced the most violations of monotonicity at lower bandwidths. Note, however, that MWTP results are not sensitive to this cut.

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

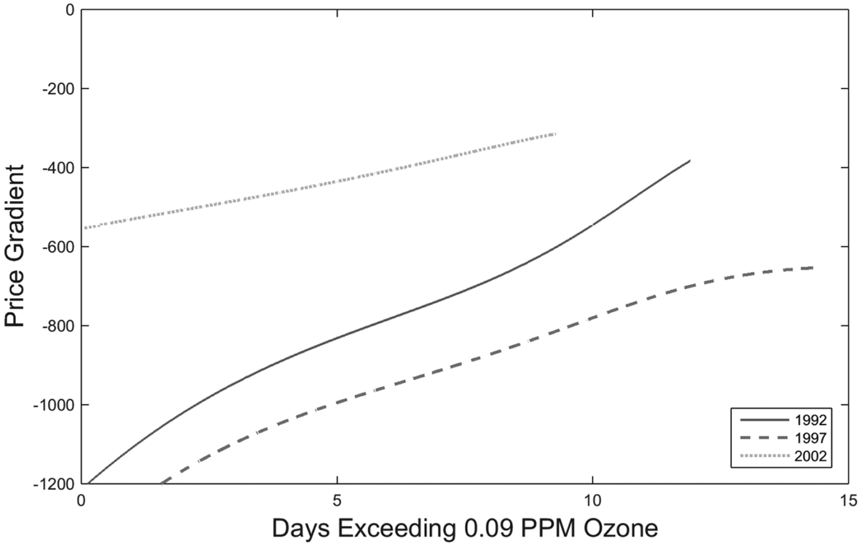

Figure 4. Hedonic price gradients (1992, 1997, 2002)

(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 (i.e. when the compensating reduction in price is larger). These features have important consequences for policy analysis and are discussed in detail in section 4.

Finally, we show the density of gradients evaluated at the point of ozone exposure in figure 5, and with 99% confidence intervals in figure A2. 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 –$801.

3.2. Results from the Second-Stage MWTP Estimation

3.2.1. Demand Estimation with the Two-Purchase Panel

Employing our approach with a panel of households that 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).

In a minority of cases (19%), the recovered slopes are positive and sufficiently steep that they exceed the positive slope of the gradients, thus violating the second-order

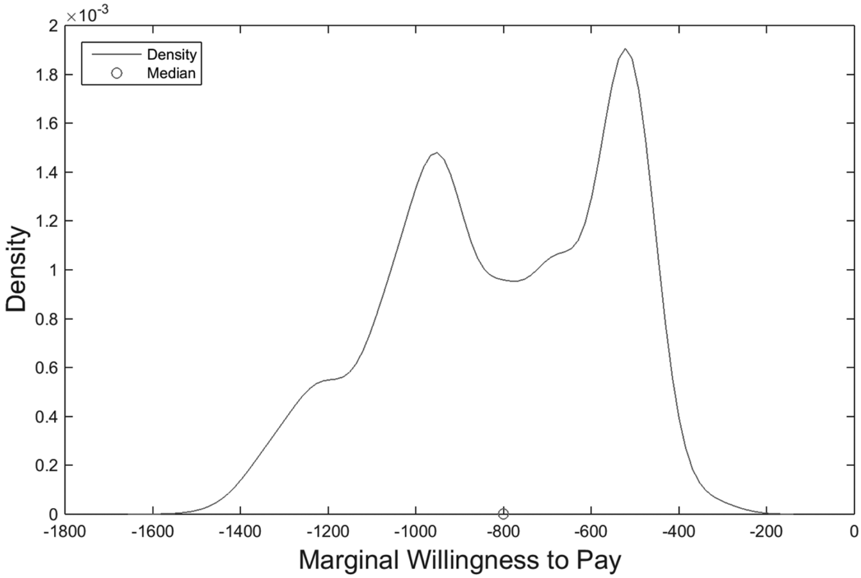

Figure 5. Distribution of local estimates of MWTP

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 a household's slope, which would explain why some households do not appear to be maximizing utility. The first is that preferences may not be stable over time (which violates our assumption of time-invariant preferences) and the second is that the merge on buyer name in the data is potentially imperfect. While our data are novel, even richer data could potentially address both of these issues. For example, richer data describing time-varying attributes would allow our MWTP curve to vary over time based on observable characteristics. Additionally, a data set 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 –\$209. However, it is difficult to summarize the joint distributions of intercepts and slopes in a single figure. Instead, we focus on the distribution of MWTP elasticities (with respect to ozone), which is calculated as

$$\hat{\alpha}_{4,i}\,\frac{Z_{j^*(i),t}}{\hat{\rho}_{j^*(i),t}}\,.$$

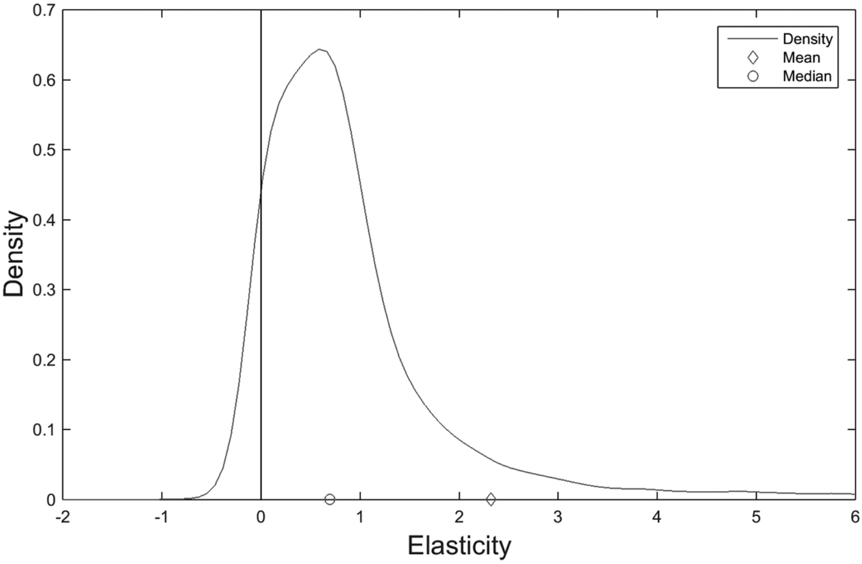

Figure 6. Distribution of MWTP elasticities

The distribution of these elasticities is shown in figure 6, and the 99% confidence intervals are shown in figure A3.

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 pollution increases, with a median value of 0.69.[31] The distribution of these elasticities also shows significant heterogeneity. While the median value is 0.69, the distribution has an interquartile range of 0.89. This elasticity is not, however, strongly correlated with the level of ozone pollution people are experiencing (i.e., a correlation of 0.01). As a partial robustness check, we also calculate the correlation between the elasticity and the gap between years of home purchases and, encouragingly, we find it to be very close to zero at 0.02.

Finally, we can decompose the intercept ($\hat{\alpha}_{3,i}$) and slope ($\hat{\alpha}_{4,i}$) coefficients according to equation (16). 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.[32] From this figure, it can be seen that Asians and whites have steeper MWTP functions

31. In a small number of cases (8%), 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 so an interior solution exists.

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

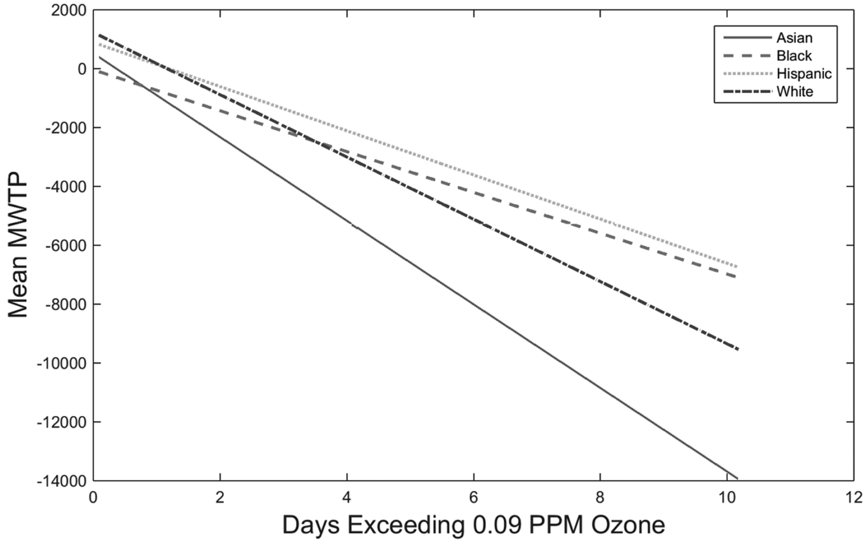

Figure 7. Average MWTP by race

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.

3.2.2. 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 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.[33] 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.

33. As noted in Bishop and Timmins (2018), 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}.

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 implement this by specifying an indirect utility function of:

$$V_{i,j,t} = \alpha_{0,i} + \alpha_{1,i}X_j + \frac{\alpha_{2,i}}{2}X_j^2 + \gamma_i \log\left(\phi - Z_{j,t}\right) + \left(I_{i,t} - R_{j,t}\right), \tag{17}$$

which yields an MWTP function of $-\gamma_i/(\phi - Z_{j,t})$. The argument of the utility function is a measure of clean air, $\phi - Z_{j,t}$, which allows the MWTP function for Z to be downward sloping.[34] The term ϕ is chosen to be the smallest number such that all values of $\phi - Z_{j,t}$ are positive, that is, $\max(Z_{j,t})$.[35] An estimate of the parameter γ_i may then be recovered as:

$$\hat{\gamma}_i = -\hat{\rho}_{j^*(i),t}\left(\phi - Z_{j^*(i),t}\right). \tag{18}$$

While this specification does allow MWTP to vary with both ozone and income, it does so in a highly restrictive way. In particular, the elasticity of MWTP with respect to clean air $(\phi - Z_{j,t})$ is restricted to be 1. Accordingly, we refer to this model as the unitary 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 elasticity model, the slope of the MWTP function (evaluated at the point of exposure) is equal to $-\gamma_i/(\phi - Z_{j,t})^2$, which given equation (18), is recovered as $\hat{\rho}_{j^*(i),t}/(\phi - Z_{j^*(i),t})$.

4. ESTIMATING THE WELFARE EFFECTS OF A NONMARGINAL CHANGE IN AIR QUALITY

To illustrate the importance of estimating heterogeneous and downward-sloping MWTP functions, we consider the valuation of a nonmarginal change in ozone exposure. Correctly estimating heterogeneous MWTP functions is critical from a policy perspective as it is important to know the correct slope of the MWTP as well as if those who are exposed to larger changes in ozone are those who are more or less sensitive to these changes. In other words, how steep are the MWTP functions and do the households exposed to the largest changes have the steeper or flatter MWTP

34. 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. Therefore, we use a measure of clean air, $\phi - Z_{j,t}$, allowing the MWTP function for Z to be negative and downward sloping.

35. In the results below, the flat MWTP model and unitary-elasticity yield similar results. If larger values of ϕ are chosen, the results become even more similar as larger values for ϕ lead to flatter MWTP slopes.

functions? These questions are particularly important given that households likely sort based on their preferences for ozone.

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. 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.[36]

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 elastic MWTP functions to highlight the importance of estimating, rather than specifying, heterogeneity in MWTP slopes.[37]

Using our model with heterogeneous slopes, we find that the mean willingness to pay to avoid a 33% increase in ozone is $1,021 and that the median willingness to pay to avoid the increase is $770. There is considerable heterogeneity as evidenced by the interquartile range of $1,106. The full distribution of welfare effects is shown in figure 8, and the 99% confidence interval is shown in figure A4.

It is well accepted in the literature that using marginal estimates to value nonmarginal changes leads to biased welfare estimates. This is effectively what is being done in the flat MWTP case.[38] In our context, this bias would lead to smaller estimated welfare effects. We find this: the welfare effects are considerably smaller under the assumption of flat MWTP curves. The mean welfare loss is only $848 and the median loss is $636.[39]

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 those from a Bajari-Benkard-style unitary elasticity model, where the MWTP function is given by $-\gamma_i/(\phi - Z_{j,t})$. The unitary elasticity model yields substantially different welfare estimates than our panel-data estimates. The mean and

36. 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.

37. 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 dropping those with very steep slopes.

38. 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.

39. The interquartile range of the distribution is $876. The full distribution of welfare effects and 99% confidence interval is shown in figure A5.

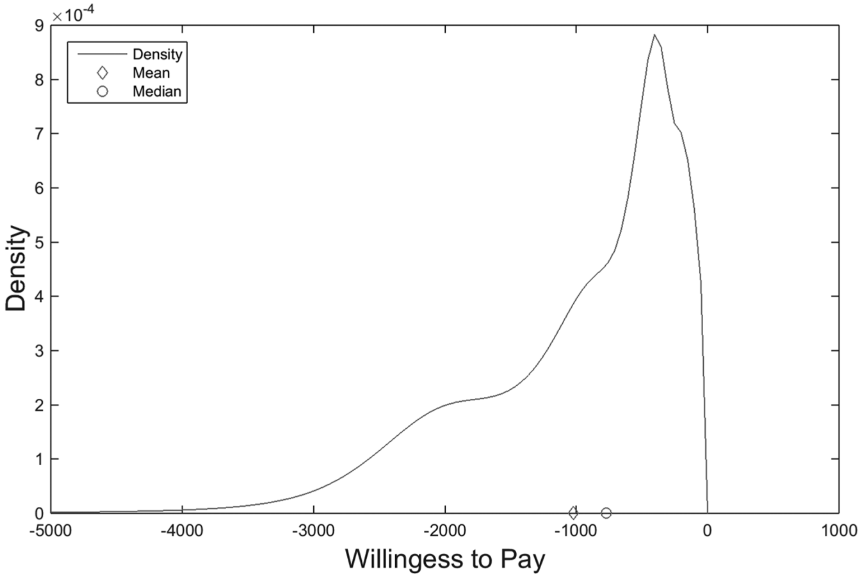

Figure 8. Welfare costs of a nonmarginal (33%) increase in ozone

median welfare costs are very close to the flat-MWTP case at $909 and $652, respectively.[40]

To understand the differences between our estimates and the unitary elasticity estimates, it is worth revisiting the distinctions between the frameworks. In both frameworks, the level of the MWTP function (at the point of consumption) is recovered by the first-order condition for utility maximization 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.[41] In the unitary 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' sensitivity to price changes is observed directly in the data: how much does the household change their ozone exposure when the implicit price of ozone changes? In this application, the MWTP function is estimated to be considerably steeper when estimated using the panel-data approach.

The unitary-elasticity model also imposes counterintuitive patterns on the heterogeneity of slopes. In the data, when the price gradient lies further from zero (i.e.,

40. The interquartile range is $930. The full distribution of welfare effects and 99% confidence interval is shown in figure A6.

41. 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.

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 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 –0.56.[42] In contrast, using panel data to estimate the slopes yields the (more sensible) result that those who are exposed to the most ozone have flatter MWTP curves on average. With our estimated slopes, the correlation between the slope of the MWTP function and the amount of ozone exposure is +0.27.

It is also worth noting that, as the policy we considered proportionally changed ozone levels, the unitary-elasticity model assigned the steepest MWTP curves to those who experienced the greatest change. If we had instead considered a policy that induced a constant change in ozone levels, the unitary-elasticity estimates would have been even closer to the flat-MWTP function case. Analogously, with a constant-change policy the panel-data estimates would have been even further from the flat-MWTP function case.

Finally, we consider a 33% reduction in ozone with results following a similar, yet opposite, pattern. In our panel-data framework with estimated slopes, the mean welfare benefit is \$675. By construction, the mean welfare benefit in the flat-MWTP framework is simply the negative of the cost for the 33% increase in ozone, that is, \$848. In the unitary elasticity framework with assigned slopes, the mean welfare benefit is again close to the flat-MWTP case at \$806.

5. 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 nonmarginal 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

42. The first factor driving this positive correlation is that $(\phi - Z)$ appears in the denominator of the MWTP function—larger values of Z therefore lead the econometrician to infer steeper MWTP curves (i.e., more negative slopes). The second factor driving this positive correlation is that $\hat{\rho}_{j^*(i),t}$ appears in the numerator of the estimate of the MWTP function. As the price gradients are upward sloping, this leads the within-year correlation between MWTP slope and ozone exposure to be positive. However, this within-year positive correlation is outweighed by a strong between-year effect which leads to negative correlation between MWTP slope and ozone.

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.

Importantly, the panel-data-based estimation approach described in this paper is increasingly implementable with the growing availability of panel data sets that match buyer attributes to the prices and characteristics of the properties they purchase. In addition, it is applicable to a wide variety of nonmarginal 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 nonmarginal policy changes without encountering the well-known endogeneity issues with the recovery of hedonic demand functions.

REFERENCES

Bajari, Patrick, and Lanier C. Benkard. 2005. Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of Political Economy* 113 (6): 1239–76.

Bajari, Patrick, and Matthew 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.

Banzhaf, H. Spencer. 2015. Panel-data hedonics: Rosen's first stage and difference-in-differences as "sufficient statistics." Unpublished manuscript, Georgia State University.

Bartik, Timothy J. 1987. The estimation of demand parameters in hedonic price models. *Journal of Political Economy* 95 (1): 81–88.

Bayer, Patrick, Robert McMillan, Alvin Murphy, and Christopher Timmins. 2016. A dynamic model of demand for houses and neighborhoods. *Econometrica* 84 (3): 893–942.

Bishop, Kelly C., and Alvin D. Murphy. 2018. Valuing time-varying attributes using the hedonic model: When is a dynamic approach necessary? *Review of Economics and Statistics* (forthcoming).

Bishop, Kelly C., and Christopher D. Timmins. 2018. Estimating the marginal willingness to pay function without instrumental variables. Unpublished manuscript, Arizona State University.

Brown, James N., and Harvey S. Rosen. 1982. On the estimation of structural hedonic price models. *Econometrica* 50 (3): 765–68.

Case, Karl E., and Robert J. Shiller. 1989. The efficiency of the market for single-family homes. *American Economic Review* 79 (1): 125–37.

Chay, Kenneth Y., and Michael Greenstone. 2005. Does air quality matter? Evidence from the housing market. *Journal of Political Economy* 113 (2): 376–424.

Clapp, John M. 2003. A semiparametric method for valuing residential location: Application to automated valuation. *Journal of Real Estate Finance and Economics* 27 (3): 303–20.

Diamond, Rebecca, and Timothy J. McQuade. 2016. Who wants affordable housing in their backyard? An equilibrium analysis of low income property development. NBER Working Paper 22204, National Bureau of Economic Research, Cambridge, MA.

Du, Pang, Christopher F. Parmeter, and Jeffrey S. Racine. 2013. Nonparametric kernel regression with multiple predictors and multiple shape constraints. *Statistica Sinica* 23:1347–71.

Ekeland, Ivar, James J. Heckman, and Lars Nesheim. 2004. Identification and estimation of hedonic models. *Journal of Political Economy* 112 (1): S60–S109.

Epple, Dennis. 1987. Hedonic prices and implicit markets: Estimating demand and supply functions for differentiated products. *Journal of Political Economy* 95 (1): 59–80.

Fan, Jianqing, and Irene Gijbels. 1996. *Local polynomial modeling and its applications*. Monographs on statistics and applied probability, no. 66. London: Chapman & Hall.

Hall, Peter, and Li-Shan Huang. 2001. Nonparametric kernel regression subject to monotonicity constraints. *Annals of Statistics* 29 (3): 624–47.

Harberger, Arnold C. 1971. Three basic postulates for applied welfare economics: An interpretive essay. *Journal of Economic Literature* 9 (3): 785–97.

Heckman, James J., Rosa L. Matzkin, and Lars Nesheim. 2010. Nonparametric identification and estimation of nonadditive hedonic models. *Econometrica* 78 (5): 1569–91.

Henderson, Daniel J., Raymond J. Carroll, and Qi Li. 2008. Nonparametric estimation and testing of fixed effects panel data models. *Journal of Econometrics* 144 (1): 257–75.

Henderson, Daniel J., and Christopher F. Parmeter. 2015. *Applied nonparametric econometrics*. Cambridge: Cambridge University Press.

Kuminoff, Nicolai V., Christopher F. Parmeter, and Jaren 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): 145–60.

Lee, Yoonseok, and Debasri Mukherjee. 2014. Nonparametric estimation of the marginal effect in fixed-effect panel data models: An application on the environmental Kuznets curve. Unpublished manuscript, Syracuse University.

Mendelsohn, Robert O. 1985. Identifying structural equations with single market data. *Review of Economics and Statistics* 67 (3): 525–29.

Murray, Michael P. 1975. The distribution of tenant benefits in public housing. *Econometrica* 43:771–88.

Parmeter, Christopher F., Daniel J. Henderson, and Subal C. Kumbhakar. 2007. Nonparametric estimation of a hedonic price function. *Journal of Applied Econometrics* 22 (3): 695–99.

Poterba, James. 1984. Tax subsidies to owner-occupied housing: An asset market approach. *Quarterly Journal of Economics* 99 (4): 729–52.

Qian, Junhui, and Le Wang. 2012. Estimating semiparametric panel data models by marginal integration. *Journal of Econometrics* 167 (2): 483–93.

Rosen, Sherwin. 1974. Hedonic prices and implicit markets: Product differentiation in perfect competition. *Journal of Political Economy* 82 (1): 34–55.

Ryan, David L., and Terence J. Wales. 2000. Imposing local concavity in the translog and generalized Leontief cost functions. *Economics Letters* 67 (3): 253–60.

Su, Liangjun, and Xun Lu. 2013. Nonparametric dynamic panel data models: Kernel estimation and specification testing. *Journal of Econometrics* 176 (2): 112–33.