A Theory-Based Approach to Hedonic Price Regressions with Time-Varying Unobserved Product Attributes: The Price of Pollution

Patrick Bajari
University of Minnesota and NBER

Jane Cooley
University of Wisconsin

Kyoo il Kim
University of Minnesota

Christopher Timmins
Duke University and NBER

October 6, 2009

Abstract

Hedonic techniques are commonly used to recover implicit prices for the attributes of differentiated products – particularly housing. However, if unobservable house or neighborhood attributes are correlated with the attributes of interest, OLS estimates of their implicit prices will be biased. A variety of quasi-experimental approaches have been used to deal with this problem in the past, but we argue that, in many important empirical settings, these strategies are not applicable. In this paper, we propose an alternative strategy based instead on two assumptions from economic theory – (1) that market prices reflect the characteristics of the home, including those that are not directly observed by the econometrician, and (2) that housing markets are informationally efficient, at least in a limited sense. In particular, we assume that characteristics observable to buyers cannot be used to earn excess returns. Using data describing housing transactions in California's Bay Area between 1990 and 2006, we find evidence in support of this assumption. Applying our estimator, we recover implicit prices for three of the EPA's "criteria" air pollutants – particulate matter (PM10), sulfur dioxide (SO2), and ground-level ozone (O3). In contrast to simple cross-sectional or fixed effect estimators, marginal willingnesses to pay for a reduction in all three pollutants (considered individually or together) are all statistically significant with the expected sign and quite large in magnitude. These results suggest that ignoring bias from time-varying correlated unobservables will lead to an understatement of the benefits of a pollution reduction policy.

1 Introduction

In a hedonic regression, the economist attempts to consistently estimate the relationship between prices and product attributes in a differentiated product market. The regression coefficients are commonly referred to as implicit prices, which can be interpreted as the effect on the market price of increasing a particular product attribute while holding the other attributes fixed. Given utility-maximizing behavior, the consumer's marginal willingness to pay for a small change in a particular attribute can be inferred from an estimate of its implicit price.

Hedonic regressions suffer from a number of well-known problems. Foremost among them, the economist is unlikely to directly observe all product characteristics that are relevant to consumers, and these omitted variables may lead to biased estimates of the implicit prices of the observed attributes. For example, in a house-price hedonic regression, the economist may observe the number of square feet, the lot size, and the average education level in the neighborhood. However, many product attributes such as curb appeal, the quality of the landscaping and the crime rate may be unobserved to the econometrician. If the unobserved attributes are correlated with the observed attributes, ordinary least squares estimates of the implicit prices will be biased. When correlated unobservables are time-invariant, they can be accounted for with fixed effects if panel data are available. When correlated unobservables vary over time, previous research has relied instead on instrumental variables, regression discontinuity, or other forms of quasi-experimental variation to avoid this bias. Chay and Greenstone (2005), Greenstone and Gallagher (2007), and Black (1999) have proposed quasi experimental approaches to this problem, exploiting either a discontinuity in the application of a regulation or a structural break due to a boundary.

If the regulation or boundary is exogenous and generates large movements in the product attributes, these methods may be attractive for estimating implicit prices for at least two reasons. First, they allow the econometrician to remove the bias from omitted variables that may confound estimates of implicit prices. Second, the identifying assumptions are transparent and the estimators are simple to implement (often using well-known statistical packages). Why would anyone choose not to adopt this straightforward approach? We argue that, in many important hedonic applications, this sort of identification cannot be achieved. First, a source of quasi-randomness that generates exogenous variation in product characteristics may not be available in a particular application. Second, even if a natural experiment is found, implicit prices may not be precisely estimated because the instruments implied by that experiment are weak. Third, there may be concern about the validity of the identifying assumptions (i.e., is there a reason to suspect that the proposed instruments may themselves be endogenous?). Fourth, a regulatory discontinuity or structural break caused by a boundary may only identify policy impacts over a narrow range, rather than over the full range of interest to policy-makers. It would therefore be useful to have an alternative set of assumptions with which to identify implicit prices. At a minimum, this alternative approach would provide a way to check the robustness of the results from a quasi-experiment; in other situations, it would provide a viable estimation strategy when quasi-random variation in the product attribute of interest cannot be assumed.

In this paper, we propose such a method. Our identifying assumptions are derived from economic theory. The first assumption is that home price in a local market at a point in time can be written as a function of a home's attributes. Importantly, for our application, we assume that this includes attributes that are observed to the home buyers but not the economist. This assumption is maintained in theoretical models underlying hedonic regressions including Rosen (1974), Epple (1987), Ekeland, Heckman and Nesheim (2004), Heckman, Matzkin and Nesheim (2003) and Bajari and Benkard (2005).

In most applications, it seems reasonable to assume that buyers have superior information about home attributes compared to the economist. For example, it is difficult, if not impossible, for the economist to directly measure the "curb appeal" of a home. However, anyone who has purchased a home knows this is an important consideration for many buyers. Our first assumption implies that curb appeal and other attributes like it are priced by the market even if the econometrician fails to measure them. As a consequence, the residual from a hedonic regression contains information that the researcher can use to price home attributes that she does not directly observe.

Our second identifying assumption is that housing markets exhibit a limited form of informational efficiency. In particular, information about a home's characteristics at the time of purchase cannot be used to make excess returns. Put differently, housing markets build into the price both current and anticipated changes in the characteristics of a home. Thus, innovations in the characteristics of a home cannot be used to earn excess returns given current information. For example, suppose that a home exhibits very high levels of ozone pollution during an abnormally hot summer (nitrogen oxides and volatile organic compounds combine with heat and sunlight to create ground-level ozone). While this high level of pollution is an unattractive feature of the house, our assumption implies that buyers rationally anticipate that when temperatures cool in subsequent years, pollution will fall and market prices will behave accordingly.

In our paper, we demonstrate that these two assumptions allow us to consistently estimate implicit prices using repeat sales data. The intuition behind our estimator is straightforward. Suppose that we observe a home that sold in 1998 and again in 2003. Our first assumption allows us to use the 1998 sales price to impute a market value for the omitted product attributes in 1998. If the market price was abnormally positive (negative) after controlling for the covariates in the econometrician's data set, we would infer that the home had a large positive (negative) value for characteristics that were not observed by the economist. More formally, the hedonic equation would allow us to form a control function to impute the market value of omitted attributes.

Our second condition provides us with a moment equation similar to well-known GMM estimators in financial econometrics. We should expect the value of omitted attributes to evolve over time. However, our limited market efficiency assumption implies that the innovation in the omitted attribute must be orthogonal to the agent's information at the time of purchase. This moment equation permits identification of the implicit prices in the hedonic equilibrium while taking into account potentially confounding unobservables. We show that our strategy will allow for a flexible functional form by casting our problem in the framework of Ai and Chen (2003). Approaches that exploit quasi-randomness may require linearity, as in the case of regression discontinuity, or frequently require a parsimonious functional form because instrumental variables do not have adequate variation to identify models with many parameters.

We admit that these two identifying assumptions are an approximation of the way housing markets function in reality. For example, our first assumption will not hold perfectly because home prices are often determined by negotiation and therefore cannot be explained exactly by the home's characteristics. At the same time, we argue that there are not many opportunities for a free lunch in a housing market with many buyers and sellers. Finding "steals", where the asking price significantly understates the value of a home's attributes, is the exception rather than the rule. Only rarely can a buyer find twice the home for half the price.

Our second assumption is also an approximation to real world housing markets. A fraction of buyers may be able to forecast which homes will have a particularly strong appreciation rate. At the same time, competition is likely to limit the returns from speculation using publicly observed information on home characteristics. Our application is a rich data set from the San Francisco Bay Area. Over the period that we study, this is a quintessential "hot" market, with a large number of buyers and sellers. Our assumption allows for the possibility that buyers anticipated profiting, perhaps handsomely, from buying and selling their homes. What our assumption rules out is that buyers were able to use commonly observed home characteristics to earn abnormal profits. Case and Shiller (1989) find support for this assumption, arguing that within a given metropolitan area, it is difficult to find variables that predict excess returns in housing. In our application, we find support for this assumption in the sense that characteristics such as square footage, lot size, and air pollution at the time of purchase cannot be used to generate economically significant excess returns in a local housing market.

As an application of our approach, we consider the value individuals place on a marginal improvement in air quality, as revealed by their home buying decisions. In particular, we analyze three of the EPA's "criteria pollutants" (i.e., pollutants used by the EPA in setting emissions regulations) – particulate matter (PM10), sulfur dioxide (SO2), and ground-level ozone (O3) – all of which are known to have adverse health consequences and impose aesthetic costs. Importantly, we expect there to be many more salient determinants of individual housing choice that our data do not describe. There is, therefore, good reason to be concerned about omitted variables bias. Moreover, if changes in pollution are correlated with changes in these omitted variables, a simple fixed effect estimator will not correct the bias.

Using data describing housing transactions in California's Bay Area between 1990 and 2006, we show evidence in support of the hypothesis that the market is informationally efficient. Using our estimator, we recover implicit prices for the three criteria air pollutants described above. In contrast to simple cross-sectional or fixed effect estimators, marginal willingnesses to pay for a reduction in all three pollutants (considered individually or together) are all statistically significant and have the expected sign. The elasticity of housing price with respect to PM10 (considered alone) is -0.30, which corresponds more closely with Chay and Greenstone's (2005) instrumental variables estimates (-0.21 to -0.35) than do other estimates in the literature. Moreover, unlike Chay and Greenstone (2005), our approach does not require that we treat the US as being comprised of a single, unified housing market. Rather, we are able to derive estimates with data from just one metropolitan area. Considered together, PM10, SO2, and O3 exhibit house price elasticities of -0.23, -0.15, and -0.26, respectively. We contrast with a simple fixed-effects model and find that controlling for time-varying unobservables appears to be extremely important for all three pollutants.

This paper proceeds as follows. Section 2 describes our estimator. Section 2.1 describes identification in the most general terms. Sections 2.2 and 2.3 describe a number of modifications and restrictions that we make to the general model before applying it to data. Section 3 describes the data that we use for our application. Section 4 presents results from our model, and compares them to results from traditional cross-sectional and fixed-effects specifications. Section 5 concludes.

2 Model: Estimating Implicit Prices

In this section, we consider the traditional hedonic framework – a model of demand in a differentiated products market in which a consumer maximizes static utility. The primary application we have in mind is housing, however, many of the methods we propose could carry over to other differentiated product markets where our assumptions are maintained. Houses, indexed by $j = 1, ..., J$, can be completely described by a finite vector of attributes. Let $\overline{x}_j$ denote a 1 by K vector of attributes such as the number of square feet, the lot size, or the year built, all of which are commonly observed by the econometrician and the consumer. In addition, let ξ_j denote a scalar that captures an omitted attribute of the house that is observed by the consumers, but not by the economist. For instance, while data sets on housing are quite detailed, they typically do not report features such as the curb appeal of a home or its state of repair, both of which may be important to buyers. For notational and expositional simplicity, we require these omitted attributes to be captured in a single product attribute, ξ_j, though many of our results allow for a more general error term with vector-valued omitted attributes. To summarize, from the perspective of a consumer $i = 1, ..., I$, product j can be completely summarized by the 1 by $(K + 1)$ vector $(\overline{x}_j, \xi_j)$.

Equilibrium prices can be written as $p_j = \mathbf{p}(\overline{x}_j, \xi_j)$. We will refer to $\mathbf{p}$ as the hedonic price function. This is a map between the product characteristics $(\overline{x}_j, \xi_j)$ and the price of good j (p_j). The hedonic price function $\mathbf{p}$ is determined in equilibrium by the interactions of buyers and sellers. Bajari and Benkard (2005) show that consumer rationality plus mild restrictions on consumer preference imply that $\mathbf{p}$ is a function, not a correspondence. As discussed in the introduction, the existence of the function $\mathbf{p}$ is our first key assumption derived from economic theory.

In empirical applications, the economist is frequently concerned with estimating $\mathbf{p}(\overline{x}_j, \xi_j)$ using data on the observed prices, p_j and characteristics, $\overline{x}_j$. Hedonic price regressions are commonly conducted assuming that $E[\xi_j | \overline{x}_j] = 0$, that is, the omitted product attributes are mean independent of the observed attributes. This assumption has been frequently criticized in the literature, going back to Small (1975). Returning to our earlier example, suppose that ξ_j reflects the curb appeal of a home. The above moment condition would imply that the expected value of curb appeal is the same for small homes in low income neighborhoods as it is for million dollar homes in exclusive neighborhoods. However, in practice we expect higher values of desirable omitted attributes to be positively correlated with higher values of desirable observed attributes. Thus, failure to correct for this omitted variable would bias upward estimates of implicit prices of desirable attributes. In empirical applications, the only proposed solutions have relied on quasi-random sources of variation such as breaks in geography (Black, 1999) or discontinuities in the application of regulations (Chay and Greenstone, 2004; Greenstone and Gallagher, 2007). While these are important contributions to the empirical literature, they may face limitations like those discussed in the introduction.

We propose an alternative approach to estimating implicit prices. We begin by considering a flexible specification, but later impose restrictions on the model to facilitate identification with the data in our Bay Area application. We consider cases in which there are data on repeat sales so that the price of home j is observed in several time periods among $t = 1, 2, ..., T$. Note that the price does not need to be observed in all time periods. Our empirical strategy will require as few as two sales for each house.

To simplify notation, consider the case where there is a single observed, potentially time-varying characteristic, $x_{j,t}$ (all results apply if this were a vector of characteristics). Suppose the system of hedonic pricing equations is:

$$\begin{aligned} \ln(p_{j,1}) &= \alpha_1 + h_1(x_{j,1}) + \xi_{j,1} \\ &\vdots \\ \ln(p_{j,T}) &= \alpha_T + h_T(x_{j,T}) + \xi_{j,T}, \end{aligned} \tag{1}$$

where we normalize $h_t(0) = 0$ and $h_t(\cdot)$ is nonparametrically specified. Since we can observe prices of homes only when they actually transact, we have an unbalanced panel where some of $\ln(p_{jt})$'s are never observed in (1).

In what follows, we assume that agents in the market are uncertain about the evolution of $\xi_{j,t}$. This uncertainty could come from one of two sources. The first is that the omitted characteristics change over time periods. For example, a noisy neighbor may move in next door to home j or an infestation may make it necessary to cut down all the large trees in home j's neighborhood. The second is that the price of the omitted attributes could change over time. In our model, we will assume that the omitted product attribute evolves according to a first order Markov process,

$$\xi_{j,t'} = \gamma(t, t')\xi_{j,t} + \eta(j, t, t'). \tag{2}$$

Here $\gamma(t, t')\xi_{j,t}$ is the expected value of the omitted attribute at time t' conditional on its value at time t, and $\eta(j, t, t')$ is the innovation in the omitted attribute.

Let I_t denote the information available to the buyer at time t. Full informational efficiency would require that "anything in the information set at time t should have no explanatory power for individual house price changes subsequent to that date." (Case and Shiller, 1989). In this case, the natural log of house price should evolve according to a random walk. We require only a limited form of informational efficiency under which the price of any home j at time periods t and t' must satisfy

$$\ln(p_{j,t\prime}) - \ln(p_{j,t}) = E[\ln(p_{j,t\prime}) - \ln(p_{j,t})|I_t] + \zeta_{j,t,t'}.$$

The left hand side variable in the above equation is the log gross return on home j between t and t', $\ln(p_{j,t\prime}) - \ln(p_{j,t})$. The right hand side is the expected value of $\ln(p_{j,t\prime}) - \ln(p_{j,t})$ given current information and the forecast error $\zeta_{j,t,t'}$. Given that innovations in $x_{j,t}$ may be forecastable, we allow forecasts of house price appreciation to depend upon them. Using equations (1) and (2), it follows that:

$$\begin{aligned}
\zeta_{j,t,t'} &= (\alpha_{t'} - \alpha_t) - E[\alpha_{t'} - \alpha_t|I_t] + \big(h_{t'}(x_{j,t'}) - h_t(x_{j,t})\big) - E[h_{t'}(x_{j,t'}) - h_t(x_{j,t})|I_t] \\
&\quad + (\xi_{j,t'} - \xi_{j,t}) - E[\xi_{j,t'} - \xi_{j,t}|I_t] \\
&= (\alpha_{t'} - \alpha_t) - E[\alpha_{t'} - \alpha_t|I_t] + \big(h_{t'}(x_{j,t'}) - h_t(x_{j,t})\big) - E[h_{t'}(x_{j,t'}) - h_t(x_{j,t})|I_t] \\
&\quad + (\gamma(t,t') - 1)\xi_{j,t} + \eta(j,t,t') - (\gamma(t,t') - 1)\xi_{j,t} \\
&= (\alpha_{t'} - \alpha_t) - E[\alpha_{t'} - \alpha_t|I_t] + \big(h_{t'}(x_{j,t'}) - h_t(x_{j,t})\big) - E[h_{t'}(x_{j,t'}) - h_t(x_{j,t})|I_t] \\
&\quad + \eta(j,t,t')
\end{aligned}$$

Our second major assumption of informational efficiency, as defined above, is equivalent to

$$E[\eta(j,t,t')|I_t] = 0. \tag{3}$$

That is, the innovation in the omitted attribute is orthogonal to the time t information set I_t. As a consequence, a rational agent's forecast of $\xi_{j,t'}$ at time t is $\gamma(t,t')\xi_{j,t}$. Combining this assumption with the above equation implies that:

$$\begin{aligned}
E[\zeta_{j,t,t'}|I_t] &= E[\alpha_{t'} - \alpha_t|I_t] - E[\alpha_{t'} - \alpha_t|I_t] + E[h_{t'}(x_{j,t'}) - h_t(x_{j,t})|I_t] \\
&\quad - E[h_{t'}(x_{j,t'}) - h_t(x_{j,t})|I_t] + E[\eta(j,t,t')|I_t] \\
&= E[\eta(j,t,t')|I_t] \\
&= 0
\end{aligned}$$

The assumption that $E[\eta(j,t,t')|I_t] = 0$ therefore implies a limited form of informational efficiency. In particular, we only assume that one cannot make excess returns by forecasting the innovation in the omitted attribute with current information, but we still allow for other ways of earning predictable returns (namely, by using forecasted values of $x_{j,t'}$ and $\xi_{j,t'}$ that are available to everyone). Full informational efficiency would rule-out any form of predictable returns.

For obvious reasons, directly testing our limited informational efficiency assumption is not possible. Later in the paper, however, we use the data from our Bay Area application to implement the Case-Shiller (1989) full informational efficiency test. In particular, we check to see if observables in the information set at time t can predict economically and statistically significant returns over the period (t, t').

2.1 Lagged Prices and Consistent Estimation of the Hedonic Equation

In this section, we rewrite our hedonic price function for period t' using information from the previous sale of house j (i.e., in period t) to eliminate $\xi_{j,t'}$. In particular, rewriting $\xi_{j,t'}$ as a function of $\xi_{j,t}$ using (2) and substituting $\ln(p_{j,t}) - \alpha_t - h_t(x_{j,t})$ for $\xi_{j,t}$, we get,

$$\begin{aligned} \ln(p_{j,t'}) &= \alpha_{t'} + h_{t'}(x_{j,t'}) + \xi_{j,t'} \\ &= \alpha_{t'} + h_{t'}(x_{j,t'}) + \gamma(t,t')\left[\ln(p_{j,t}) - \alpha_t - h_t(x_{j,t})\right] + \eta(j,t,t') \\ &= \left(\alpha_{t'} - \gamma(t,t')\alpha_t\right) + \gamma(t,t')\ln(p_{j,t}) \\ &\quad -\gamma(t,t')h_t(x_{j,t}) + h_{t'}(x_{j,t'}) + \eta(j,t,t'). \end{aligned} \tag{4}$$

We note that $x_{j,t'}$ could be correlated with $\eta(j,t,t')$, for example, the innovation in "curb appeal" between t and t' might have been correlated with an observable characteristic such as test scores in local public schools. This means a regression based on (4) will produce inconsistent estimates of the hedonic price function. However, by exploiting the process that describes the evolution of $x_{j,t}$ over time, we can still obtain consistent estimates for all the parameters in (4). We assume

$$x_{j,t'} = g_{t,t'}(x_{j,t}, w_{j,t}) + v_{j,t,t'} \qquad E[v_{j,t,t'}|I_t] = 0 \tag{5}$$

In words, the innovation in the observed attributes is orthogonal to time t information where $w_{j,t}$ denotes other observable variables in I_t. Also we assume that

$$\eta(j,t,t') = \tau(t,t')v_{j,t,t'} + \varepsilon_{j,t,t'}. \tag{6}$$

These assumptions imply that $x_{j,t'}$ evolves according to the process described in (5), but that the innovation in $x_{j,t'}$ may be correlated with the innovation in the omitted attribute. Applying assumptions (3) and (6) to (4), we obtain

$$\begin{aligned}\ln(p_{j,t'}) &= \left(\alpha_{t'} - \gamma(t,t')\alpha_t\right) + \gamma(t,t')\ln(p_{j,t}) \\ &\quad -\gamma(t,t')h_t(x_{j,t}) + h_{t'}(x_{j,t'}) + \tau(t,t')v_{j,t,t'} + \varepsilon_{j,t,t'}.\end{aligned}$$

Our identification and estimation methods are then based on the following moment condition

$$E\left[\varepsilon_{j,t,t'}|I_t, v_{j,t,t'}\right] = 0.$$

This moment condition states that after controlling for time t information I_t and the innovation in the observed attributes $v_{j,t,t'}$, the innovation in the omitted attribute has an expected value of zero. This moment condition motivates the use of a control function approach – in the first step, estimate equation (5) and obtain fitted residuals, $\widehat{v}_{j,t,t'}$; in the second step, include $\widehat{v}_{j,t,t'}$ as an additional regressor in (4). This procedure differs considerably from standard methods for estimating hedonics where the identifying assumption for consistent estimation is that $E\left[\xi_j|\overline{x}_j\right] = 0$.

Intuitively, our approach uses the information in lagged prices, $p_{j,t}$ to impute a lagged value of the omitted attribute. For example, if the price for home j was unusually high after controlling for $\overline{x}_j$, we would infer that $\xi_j = \ln(p_{j,t}) - \alpha_t - h_t(x_{j,t})$ was also large. This is where our first economic assumption has bite. We assume that prices reflect attributes that are observed to consumers, but not to the econometrician. We assume that the innovations ($v_{j,t,t'}$) in the observed attributes are orthogonal to current information. As in the previous section, if this were not true, it would be possible to earn excess returns in the housing market. We also assume that after controlling for I_t and $v_{j,t,t'}$, the innovation in the omitted attribute $\varepsilon_{j,t,t'}$, is mean zero. As we discussed in the last section, this is necessary to prevent buyers from earning excess returns in the housing market given time t information. Putting it all together, we achieve identification through our two economic assumptions – (i) prices reflect product attributes, including those that are unobserved by the economist, and (ii) buyers cannot use current information to make excess returns in housing. We develop this more formally in the next section.

2.2 Identification and Estimation

To consider a general setting, suppose for each house j we observe transaction prices on three occasions, denoted by $t_a(j)$, $t_b(j)$, and $t_c(j)$ such that $1 \leq t_a(j) < t_b(j) < t_c(j) \leq T$ (we later consider in detail the case when we observe only two transactions). For now, we also assume that

$x_{j,t}$ is time-varying, but later relax this assumption.

We write (4) for several time periods (assuming that we have enough observations of transaction prices for each time period[1]) such that

$$\begin{aligned}
\ln(p_{j,t_b}) &= \alpha_{t_b} - \gamma(t_a,t_b)\alpha_{t_a} + \gamma(t_a,t_b)\ln(p_{j,t_a}) - \gamma(t_a,t_b)h_{t_a}(x_{j,t_a}) \\
&\quad + h_{t_b}(x_{j,t_b}) + \tau(t_a,t_b)\nu_{j,t_a,t_b} + \varepsilon_{j,t_a,t_b} \\
\ln(p_{j,t_c}) &= \alpha_{t_c} - \gamma(t_b,t_c)\alpha_{t_b} + \gamma(t_b,t_c)\ln(p_{j,t_b}) - \gamma(t_b,t_c)h_{t_b}(x_{j,t_b}) \\
&\quad + h_{t_c}(x_{j,t_c}) + \tau(t_b,t_c)\nu_{j,t_b,t_c} + \varepsilon_{j,t_b,t_c}.
\end{aligned}$$

Estimation proceeds based on the following moment conditions:

$$E[x_{j,t_b} - g_{t_a,t_b}(x_{j,t_a}, w_{j,t_a})|1, x_{j,t_a}, w_{j,t_a}] = 0 \quad (7)$$

$$E[\varepsilon_{j,t_a,t_b}|1, \ln(p_{j,t_a}), x_{j,t_a}, x_{j,t_b}, v_{j,t_a,t_b}, w_{j,t_a}] = 0 \quad (8)$$

$$E[x_{j,t_c} - g_{t_b,t_c}(x_{j,t_b}, w_{j,t_b})|1, x_{j,t_b}, w_{j,t_b}] = 0 \quad (9)$$

$$E[\varepsilon_{j,t_b,t_c}|1, \ln(p_{j,t_b}), x_{j,t_b}, x_{j,t_c}, v_{j,t_b,t_c}, w_{j,t_b}] = 0 \quad (10)$$

where $w_{j,t}$ denotes other observable covariates in I_t. From (7) we identify $g_{t_a,t_b}(\cdot)$ (along with v_{j,t_a,t_b}), and from (8) we identify α_{t_a}, α_{t_b}, $\gamma(t_a,t_b)$, $h_{t_a}(x_{j,t_a})$, $h_{t_b}(x_{j,t_b})$, and $\tau(t_a,t_b)$. Similarly from (9) we identify $g_{t_b,t_c}(\cdot)$ (along with v_{j,t_b,t_c}) and from (10) we identify α_{t_b}, α_{t_c}, $\gamma(t_b,t_c)$, $h_{t_b}(x_{j,t_b})$, $h_{t_c}(x_{j,t_c})$, and $\tau(t_b,t_c)$.

These may be run as two separate sets of estimations – i.e., one is based on (7) and (8) and the other based on (9) and (10). We note, however, that α_{t_b} and $h_{t_b}(x_{j,t_b})$ are over-identified from the moment conditions, which suggests combining all the moment conditions and performing a nonlinear nonparametric estimation. Having set-up the moment conditions in (7)-(10), one can cast them into Ai and Chen (2003)'s framework and estimate all the parameters (including nonparametric functions) simultaneously. The consistency of the estimators and the asymptotic normality of the parametric components can be obtained following Ai and Chen (2003).

Once we estimate the hedonic function, another parameter of interest will be the weighted average derivative of the log housing price ($\ln(p_{j,t'})$) with respect to the observed characteristic $x_{j,t}$ for $t \leq t'$, defined by

$$E\left[\lambda(x_{j,t})\frac{\partial \ln(p_{j,t'})}{\partial x_{j,t}}\right]. \quad (11)$$

[1] To be precise, $\frac{1}{J}\sum_{j=1}^{J} 1\{t_a(j) = t \text{ or } t_b(j) = t \text{ or } t_c(j) = t\} \underset{J\to\infty}{\longrightarrow} C > 0$ for all $t = 1, \ldots, T$.

This can be estimated using its sample analogue,

$$\frac{1}{J}\sum_{j=1}^{J}\lambda(x_{j,t})\frac{\partial\ln(p_{j,t'})}{\partial x_{j,t}} \approx \frac{1}{J}\sum_{j=1}^{J}\lambda(x_{j,t})\frac{\partial\widehat{\ln(p_{j,t'})}}{\partial x_{j,t}}$$

where $\widehat{\ln(p_{j,t'})}$ is the fitted log price function. The weight $\lambda(\cdot)$ satisfies $\lambda(\cdot) \geq 0$ and $\int \lambda(x)dx = 1$. In the literature, estimated implicit prices $\frac{\partial\widehat{\ln(p_{j,t'})}}{\partial x_{j,t}}$ are commonly used to recover valuation for non-market amenities such as clean air or public school quality. Since $\frac{\partial \ln(p_{j,t'})}{\partial x_{j,t}}$ varies across homeowners in the population, the function $\lambda(x_{j,t})$ allows the researcher to aggregate these heterogeneous marginal benefits. The scalar (11) summarizes the average marginal benefit from an increase in the amount of some characteristic $x_{j,t}$ in time period t and is often used to measure welfare from a policy change.

2.3 Time-invariant Covariates and Model Restrictions

When $x_{j,t}$ has no time-varying components, some of the parameters in the full model are not identified without further restrictions. Replacing $x_{j,t}$ with the time invariant covariate z_j, equation (4) becomes

$$\ln(p_{j,t'}) = \big(\alpha_{t'} - \gamma(t,t')\alpha_t\big) + \gamma(t,t')\ln(p_{j,t}) - \gamma(t,t')h_t(z_j) + h_{t'}(z_j) + \eta(j,t,t').$$

We cannot therefore identify $h_{t'}(z_j)$ separately from $h_t(z_j)$ – i.e., a multicollinearity problem. Restricting $\alpha_t = \alpha_0 \; \forall \; t$, the above equation becomes

$$\ln(p_{j,t'}) = \alpha_0\left(1 - \gamma(t,t')\right) + \gamma(t,t')\ln(p_{j,t}) - \gamma(t,t')h_t(z_j) + h_{t'}(z_j) + \eta(j,t,t').$$

With these restrictions, we can identify $\gamma(t,t')$ from the coefficient on $\ln(p_{j,t})$ and α_0 from the constant term. By further assuming $h_{t'}(z_j) = h_t(z_j)$, that function is also identified. Alternatively, one can normalize $h_1(z_j) = 1$. Then $h_t(z_j)$, $t > 1$ is identified recursively up to this normalization using the fact that $-\gamma(t,t')h_t(z_j) + h_{t'}(z_j)$ can be recovered in each period.

Imposing some structure on $\gamma(t,t')$ yields a set of over-identifying restrictions. For example, we can let $\gamma(t,t') = \gamma(t,\widetilde{t})\gamma(\widetilde{t},t')$ for $\widetilde{t}$ between t and t'. Even when some elements of x are time varying but the variation is not sufficient, one can still impose the above restrictions to obtain more stable and robust estimates.

2.4 Simple Parametric Model with Two Transactions

Consider the case when only two transactions per house are available in the data, as in our application. Moreover, impose simple parametric functional forms on the process governing the evolution of $x_{j,t}$ and on the way that variable enters the hedonic price function – i.e., $h_t(x_{j,t}) = \beta x_{j,t} \ \forall \ t$. The model simplifies to

$$\begin{aligned} v_{j,t_a,t_b} &= x_{j,t_b} - \pi_0(t_a,t_b) - \pi_1(t_a,t_b)x_{j,t_a} - \pi_2(t_a,t_b)w_{j,t_a} \\ \ln(p_{j,t_b}) &= \alpha_{t_b} - \gamma(t_a,t_b)\alpha_{t_a} + \gamma(t_a,t_b)\ln(p_{j,t_a}) + \gamma(t_a,t_b)\beta x_{j,t_a} \\ &\quad + \beta x_{j,t_b} + \tau(t_a,t_b)v_{j,t_a,t_b} + \varepsilon_{j,t_a,t_b} \end{aligned} \tag{12}$$

where t_a denotes the time period of the first sale and t_b denotes the time period of the second sale. Clearly all the parameters are identified. More importantly, the estimation can proceed as a simple application of classical two-step nonlinear least squares (2SLS) because we can rewrite (12) as

$$\begin{aligned} x_{j,t_b} &= \pi_0(t_a,t_b) + \pi_1(t_a,t_b)x_{j,t_a} + \pi_2(t_a,t_b)w_{j,t_a} + v_{j,t_a,t_b} \\ \ln(p_{j,t_b}) &= \alpha_{t_b} - \gamma(t_a,t_b)\alpha_{t_a} + \gamma(t_a,t_b)\ln(p_{j,t_a}) + \gamma(t_a,t_b)\beta x_{j,t_a} \\ &\quad + \beta\left(\pi_0(t_a,t_b) + \pi_1(t_a,t_b)x_{j,t_a} + \pi_2(t_a,t_b)w_{j,t_a}\right) + u_{j,t_a,t_b} \end{aligned} \tag{13}$$

where $u_{j,t_a,t_b} = \beta v_{j,t_a,t_b} + \eta(j,t_a,t_b)$ and $E[u_{j,t_a,t_b}|1,\ln(p_{j,t_a}),x_{j,t_a},w_{j,t_a}] = 0$. Note that this 2SLS approach is more robust to possible misspecification of (6) because (13) holds regardless of (6).

The coefficients in the second step equation (13) are nonlinear functions of $\theta(t_a,t_b) = (\alpha_{t_a},\alpha_{t_b},\gamma(t_a,t_b),\beta)'$ after we estimate $\pi_0(t_a,t_b)$, $\pi_1(t_a,t_b)$, and $\pi_2(t_a,t_b)$ in the first step. We obtain estimates by solving

$$\begin{aligned} \widehat{\pi}(t_a,t_b) &= \text{argmin}_{\pi(t_a,t_b)} \sum\nolimits_{j=1}^{J} \left\{x_{j,t_b} - \pi_0(t_a,t_b) - \pi_1(t_a,t_b)x_{j,t_a} - \pi_2(t_a,t_b)w_{j,t_a}\right\}^2 \\ \widehat{\theta}(t_a,t_b) &= \text{argmin}_{\theta(t_a,t_b)} \sum\nolimits_{j=1}^{J} \left\{\ln(p_{j,t_b}) - g(\ln(p_{j,t_a}),x_{j,t_a},\widehat{x}_{j,t_b};\theta(t_a,t_b))\right\}^2 \end{aligned}$$

where $\widehat{x}_{j,t_a,t_b} = \widehat{\pi}_0(t_a,t_b) + \widehat{\pi}_1(t_a,t_b)x_{j,t_a} + \widehat{\pi}_2(t_a,t_b)w_{j,t_a}$ and

$$\begin{aligned} g(\ln(p_{j,t_a}),x_{j,t_a},\widehat{x}_{j,t_b};\theta(t_a,t_b)) = \alpha_{t_b} - \gamma(t_a,t_b)\alpha_{t_a} + \gamma(t_a,t_b)\ln(p_{j,t_a}) \\ + \gamma(t_a,t_b)\beta x_{j,t_a} + \beta\widehat{x}_{j,t_b} + \varepsilon_{j,t_a,t_b}. \end{aligned}$$

We can also impose some parametric restrictions on $\pi(t_a,t_b)$'s and $\gamma(t_a,t_b)$'s in the above.

Note that the first step estimation contributes to the asymptotic variance of the second step estimators. We can obtain correct standard errors by following Murphy and Topel (1985). Denote

$$
\begin{aligned}
\sqrt{J}(\widehat{\pi}(t_a,t_b)-\pi(t_a,t_b)) &\rightarrow {}_d N(0,V(t_a,t_b)) \\
G(\cdot;\theta(t_a,t_b)) &= \frac{\partial}{\partial\theta(t_a,t_b)} g(\cdot;\theta(t_a,t_b)) \\
\Omega_0(t_a,t_b) &= E[\eta^2(\cdot,t_a,t_b)G(\cdot;\theta(t_a,t_b))G(\cdot;\theta(t_a,t_b))'] \\
Q_0(t_a,t_b) &= E\left[G(\cdot;\theta(t_a,t_b))G(\cdot;\theta(t_a,t_b))'\right] \\
Q_1(t_a,t_b) &= E\left[G(\cdot;\theta(t_a,t_b))\beta(1,x_{j,t_a},w_{j,t_a})\right]
\end{aligned} \tag{14}
$$

Then, we have

$$\sqrt{J}(\widehat{\theta}(t_a,t_b)-\theta(t_a,t_b)) \rightarrow_d N(0,\Sigma(t_a,t_b))$$

where

$$\Sigma(t_a,t_b) = Q_0(t_a,t_b)^{-1}\left[\Omega_0(t_a,t_b)+Q_1(t_a,t_b)V(t_a,t_b)Q_1(t_a,t_b)'\right]Q_0(t_a,t_b)^{-1}.$$

A consistent estimator of the heteroskedasticity robust variance matrix $\Sigma(t_a,t_b)$ is obtained using the following sample counterparts of (14):

$$
\begin{aligned}
\widehat{V}(t_a,t_b) &= \begin{array}{c}\left(\frac{1}{J}\sum_{j=1}^{J}(1,x_{j,t_a},w_{j,t_a})'(1,x_{j,t_a},w_{j,t_a})\right)^{-1}\left(\frac{1}{J}\sum_{j=1}^{J}\widehat{v}^2_{j,t_a,t_b}(1,x_{j,t_a},w_{j,t_a})'(1,x_{j,t_a},w_{j,t_a})\right) \\ \times\left(\frac{1}{J}\sum_{j=1}^{J}(1,x_{j,t_a},w_{j,t_a})'(1,x_{j,t_a},w_{j,t_a})\right)^{-1}\end{array}, \\
\widehat{\Omega}_0(t_a,t_b) &= \frac{1}{J}\sum_{j=1}^{J}\widehat{\eta}^2(j,t_a,t_b)G(\cdot;\widehat{\theta}(t_a,t_b))G(\cdot;\widehat{\theta}(t_a,t_b))', \\
\widehat{\eta}(j,t_a,t_b) &= \ln(p_{j,t_b})-g(\ln(p_{j,t_a}),x_{j,t_a},x_{j,t_b};\widehat{\theta}(t_a,t_b)), \\
\widehat{Q}_0(t_a,t_b) &= \frac{1}{J}\sum_{j=1}^{J}G(\cdot;\widehat{\theta}(t_a,t_b))G(\cdot;\widehat{\theta}(t_a,t_b))', \\
\widehat{Q}_1(t_a,t_b) &= \frac{1}{J}\sum_{j=1}^{J}G(\cdot;\widehat{\theta}(t_a,t_b))\widehat{\beta}(1,x_{j,t_a},w_{j,t_a}), \text{ and} \\
\widehat{\Sigma}(t_a,t_b) &= \widehat{Q}_0(t_a,t_b)^{-1}\left[\widehat{\Omega}_0(t_a,t_b)+\widehat{Q}_1(t_a,t_b)\widehat{V}(t_a,t_b)\widehat{Q}_1(t_a,t_b)'\right]\widehat{Q}_0(t_a,t_b)^{-1}.
\end{aligned}
$$

3 Data

We demonstrate the role of efficient housing markets in controlling for time-varying, correlated unobservables by measuring the marginal willingness to pay to avoid exposure to three of the

EPA's "criteria" air pollutants - particulate matter (PM10), sulfur dioxide (SO2), and ground-level ozone (O3).[2] Without extremely detailed data describing the evolution of neighborhood attributes, correlated unobservables are likely to play an important role in such an application.

We consider housing transactions from California's Bay Area (specifically, Alameda, Contra Costa, Marin, San Francisco, San Mateo, and Santa Clara counties) over the period 1990-2006. These data were purchased from the DataQuick Corporation and contain information describing the universe of housing transactions (i.e., buyers', sellers' and lenders' names, dates, loan amounts, and transaction prices) and the houses that transacted (i.e., square footage, lot size, year built, number of rooms, and how many of those rooms are bedrooms or bathrooms). Important for our purposes, the data also provide the exact street address of each home, with which we can impute pollution measures using data from thirty-seven monitors located throughout the Bay Area.

3.1 Housing Data

DataQuick reports a house's attributes as they were measured at the time of the last sale entered in our data. Because houses may have been altered (either improved or suffered some severe damage), these attributes may not be applicable to all observed transactions. We therefore carry-out a number of data cuts to avoid this problem. First, we consider the appreciation rate exhibited by each house over each pair of sales that we observe in the data. From this, we deduct the average appreciation rate for all houses that sold in the same pair of years. We then drop the houses in the top and bottom 10% of the resulting distribution of normalized appreciation rates. As such, we eliminate any house that appreciated or depreciated at a very high rate relative to other houses on the market at the same time.

Second, we drop problematic observations – for example, all observations where "year built" is missing, or where "year built" comes after the transaction date (signaling a purchase of land on which a house was then constructed). We also drop all properties that fail to report a transaction price or a latitude and longitude, houses with outlier attributes, and all observations with housing attributes that appear to be coded with error - in particular, houses where the number of bedrooms or bathrooms is greater than five. We also drop any house more than 5,000 square feet in size, or which sits on more than a 70,000 square-foot lot. We finally drop all homes that sell more than

[2] The list of criteria pollutants also includes nitrogen oxides, lead, and carbon monoxide. This list forms the basis for the EPA's primary (health) and secondary (environmental and aesthetic) emissions reduction targets. Of the six criteria pollutants, particulate matter and ground-level ozone are commonly considered to pose the greatest health threat. (http://www.epa.gov/air/urbanair)

two times in the seventeen year period we are considering. This is done primarily for the sake of convenience, as it allows us to implement our estimator using a simple specification that is easy to describe. In the end, these cuts leave us with data describing repeat transactions for 74,892 unique housing units. Table 1 summarizes the attributes of these houses.

Table 1: House Attributes (N=74,892)

	Mean	Std Dev	Minimum	Maximum
Lot size	6,983	5,799	1,000	69,900
Square feet	1720	642	500	5,000
No. of bathrooms	1.990	0.6536	1	5
No. of bedrooms	3.236	0.8258	1	5
No. of rooms	6.828	2.279	0	110
Year built	1967	22.16	1873	2005

Figure 1 describes the median transaction price in each year of our data. This makes clear that there were periods of (slow) depreciation and (rapid) appreciation in the Bay Area over the period we are considering.

Figure 1: Median Transaction Price by Year With 25th and 75th Percentiles

3.2 Air Quality Data

We measure individuals' average marginal willingness-to-pay (MWTP) to avoid three of the EPA's major criteria air pollutants.[3] The MWTP is a key determinant of the benefits of any new air pollution regulation, such as the Clean Air Act Amendments of 1990 that allowed for trading in permits to emit sulfur dioxide. The other main source of value from a new air pollution regulation comes from avoided mortality; this is typically measured by ascribing the value of a statistical life (VSL) to each death avoided by the policy.

We first consider PM10, which denotes particles less than ten micrometers in diameter. These particles (especially those smaller than 2.5 micrometers) can travel deep into the lungs and even into the bloodstream. This can lead to a variety of health problems, including asthma, chronic bronchitis, and heart attack.[4] Fine particles also reduce visibility, and prolonged exposure to PM10 can damage structures and stain building materials. While not necessarily as important as health effects from a welfare perspective, these aesthetic effects may have a marked impact on housing prices. We consider the average annual PM10 concentration, which is measured in micrograms per cubic meter (μg/m3). PM10 concentration at each house is imputed with an inverse-squared-distance weighted average of the concentrations measured at each of the thirty-seven monitoring stations in the Bay Area.

Our second pollutant is sulfur dioxide (SO2). The primary health consequences of sulfur dioxide come in the form of breathing difficulties, especially for those who suffer from asthma. Like PM10, SO2 can also create haze that impairs visibility. Acid rain (or acid fog), which is produced when SO2 reacts with water and other chemicals in the air, will damage building materials and kill vegetation. SO2 (and the remainder of our pollutants) is measured in parts per million (ppm), and we use the maximum one-hour observation observed over the course of the year at each monitor (imputed for each house again using an inverse-squared-distance weighted average of all monitors' observations). The maximum one-hour observation is an important figure used by the California Air Resources Board in determining whether or not an air district is in compliance with state regulations.

[3]Information on the health and aesthetic costs of each of the pollutants discussed in this section can be found at the EPA's web-site (`http://www.epa.gov/air/`).

[4]The Harvard "Six City" Study (Dockery et al., 1993) established many of these effects, which have been confirmed by numerous studies since that time. Lin et al., 2002; Norris et al., 1999; Slaughter et al., 2003; and Tolbert et al., 2000) have demonstrated detrimental effects, particularly for the young and elderly suffering from asthma. Hong et al., 2002; Tsai et al., 2003, and D'Ippoliti et al., 2003 provide evidence of increased risk of heart attack and stroke. Ghio et al. 2000 finds evidence of lung tissue imflammation, while Pope et al., 2002 finds increased risk of lung cancer. More recently, Samet et al., 2004 has found evidence of increased risk of heritable diseases from exposure to fine particulates.

Third, we consider ground-level ozone (O3). Similar to smog, ozone can cause a variety of severe respiratory problems including coughing, wheezing, breathing pain, aggravated asthma, and increased susceptibility to bronchitis. Exposure to peak concentrations of ground-level ozone can have acute effects, and repeated exposure to even moderate levels can lead to permanent lung damage. In addition to its health consequences, O3 has detrimental impacts on the growth of vegetation (particularly trees and other plants in urban settings), which can have important aesthetic consequences for housing prices.

Figure 2 describes the time path of three pollution measures (along with nitrogen oxides) over the sample period. To make the numbers more easily interpretable on the same graph, we express PM10 pollution in (μg/m3)*(1/1000).

Figure 2: Median One Hour Maximum Pollution Concentrations

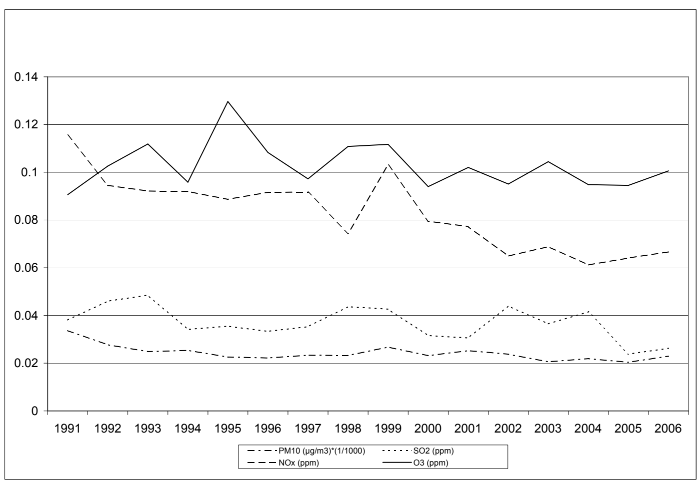

Table 2 describes the correlations across all three pollutants observed at the time of every transaction in our sample. Collinearity is not a particularly important problem in the pollutants that we consider. We therefore estimate a model with all three pollutants appearing simultaneously. In addition, we also measure the MWTP for each pollutant considered one at a time.

A final feature of these pollutants that we do not deal with is the fact that the disutility from each may be a complicated nonlinear function of the concentrations of all the other pollutants. This is a result of the photochemical processes through which they interact. See Muller, Tong,

Table 2: Correlations of Pollutants

	PM10	SO2	O3
PM10	1.0000		
SO2	0.1210	1.0000	
O3	0.2300	0.0948	1.0000

and Mendelsohn (2008) for an example of research that considers these interactions.

4 Results

4.1 Testing the Efficient Housing Market Hypothesis

We provide empirical evidence in support of our limited efficient housing market assumption by approximating Case and Shiller's (1989) test of full informational efficiency. In particular, we check to see whether or not anything in the information set at time t has any explanatory power for price changes after that time. This is a stronger test than we require. In particular, our estimator allows for predictable changes in observables and unobservables based on their current values to influence price changes. However, if we find that this predictive power (of current observables) is weak, statistically and/or economically, it bolsters our assumption that $E[\eta(j,t,t')|I_t] = 0$.

Begin by letting t_a, t_b, and t_c denote the times at which sales are observed for houses in our data set that transact three times, $t_a < t_b < t_c$. a, b, and c can be different for each of these houses. We regress the annualized return, $\frac{\ln(p_{j,t_c})-\ln(p_{j,t_b})}{t_c-t_b}$ on the average return of previous sales (which is allowed to differ across counties and years), housing attributes, pollutants, and county fixed effects.[5] While many of the coefficients are statistically significant, their economic magnitudes in terms of marginal returns are negligible. Table 3 summarizes the results.

To give a sense of the extent to which knowing housing attributes can help to generate excess returns, we calculate dollar amounts of excess returns from the estimation results in Table 3 under the scenario that the previous average return is higher by 10 percent, the lot size is larger by 100 square feet, the home size is larger by 100 square feet, the number of bathrooms is larger by 1, the number of bedrooms is larger by 1, and the three pollutant measures PM10, SO2, and O3 are

[5]The average return from previous sales is obtained as the fitted values from the regression of $\ln(p_{j,t_b}) - \ln(p_{j,t_a})$ on dummies indicating the year of the first and the second sales and the county in which the house is located.

Table 3: Efficient Market Hypothesis Testing (N=16656)

	Avg. return	Lotsize	Sqft	Bathroom	Bedroom	PM10	SO2	O3
Coeff	0.0205	0.0000	-0.0000	-0.0024	0.0077	-0.0040	1.0785	0.2496
t-stat	2.35	1.29	-13.66	-1.67	8.56	-4.80	6.78	1.03

t-statistics are calculated from clustered robust standard errors, clustered by county. The dependent variable is the annualized return.

higher by 2 $\mu g/m^3$, 5 ppb, 10 ppb, respectively. We also assume that the home prices at the time of purchase are 0.4 million, 1 million, and 2 million dollars, respectively. The results are reported in Table 4. We see that the amounts are small compared to the home prices (i.e., less than 1%). For example, when the home price is 1 million at the time of purchase, one would make excess returns of 33 dollars per year by purchasing a home with its lot size larger by 100 square feet and would earn 7,699 dollars of additional annual returns by purchasing a home with five bedrooms instead of four bedrooms. One can interpret calculations of excess returns in other cases similarly.

These results suggest that information available at time t cannot be used to affect, in an economically significant way, the returns derived from a house purchase decision. This implies that information available at that time is not particularly useful in predicting $\eta(j,t,t')$; hence, supporting our assumption that $E[\eta(j,t,t')|I_t] = 0$.

Table 4: Excess returns in dollar amounts

Home with	Avg. ret 10% ↑	Lotsize 100 sf ↑	Sqft 100 sf ↑	# Bath 1 ↑	# Bed 1 ↑	PM10 2 $\mu g/m^3$ ↑	SO2 5 ppb ↑	O3 10 ppb ↑
Price at purchase	Excess returns per year							
0.4M	820	13	-916	-953	3,080	-3,241	2,157	998
1M	2,050	33	-2,290	-2,382	7,699	-8,103	5,393	2,495
2M	4,100	66	-4,580	-4,764	15,398	-16,206	10,785	4,991

These values are calculated based on estimates in Table 3.

4.2 Measuring the Marginal Willingness-to-Pay to Avoid Air Pollution

In our application, we allow Bay Area housing prices to be determined by different hedonic price functions in each of three separate periods: (1) 1990-1994, (2) 1995-2000, and (3) 2001-2006. These periods correspond (roughly) to periods of depreciation, appreciation, and very rapid appreciation in this housing market. We report results for three different econometric models. First, we estimate a simple cross-sectional model for each of the three time periods in our data set. This approach does nothing to control for unobservables (time-varying or time-invariant) that may be correlated with pollution:

Cross-Sectional Model

$$\begin{aligned}\ln(p_{j,1}) &= \alpha_1 + x'_{j,1}\beta_1 + z'_j\phi_1 + \xi_{j,1},\\ \ln(p_{j,2}) &= \alpha_2 + x'_{j,2}\beta_2 + z'_j\phi_2 + \xi_{j,2},\\ \ln(p_{j,3}) &= \alpha_3 + x'_{j,3}\beta_3 + z'_j\phi_3 + \xi_{j,3}.\end{aligned}$$

Second, we estimate a house fixed-effect model that constrains the derivative of $\ln(P)$ with respect to each pollutant to be constant over time. This constraint allows us to recover an implicit price for pollution from the fixed-effect specification. We allow the marginal effects of other housing attributes to vary over time, meaning that we can only recover the change in the implicit prices of these attributes. The house fixed-effect model uses panel data to control non-parametrically for any time-invariant unobservables that might be correlated with pollution.

House Fixed-Effect Model

$$\begin{aligned}\ln(p_{j,3}) - \ln(p_{j,2}) &= \rho_{2,3} + (x_{j,3} - x_{j,2})'\beta + z'_j\chi_{2,3} + u_{j,2,3},\\ \ln(p_{j,3}) - \ln(p_{j,1}) &= \rho_{1,3} + (x_{j,3} - x_{j,1})'\beta + z'_j\chi_{1,3} + u_{j,1,3},\\ \ln(p_{j,2}) - \ln(p_{j,1}) &= \rho_{1,2} + (x_{j,2} - x_{j,1})'\beta + z'_j\chi_{1,2} + u_{j,1,2}.\end{aligned}$$

where $\rho_{t,t'} = (\alpha_{t'} - \alpha_t)$ and $\chi_{t,t'} = (\phi_{t'} - \phi_t)$.

Finally, we estimate a constrained specification of the model described in equation (4). We restrict $\gamma(1,3) = \gamma(1,2)\gamma(2,3)$, and $h_t(x_{j,t}) = x'_{j,t}\beta$. Constraining the marginal effect of pollution on price to be constant over time assists with model identification and makes the results more directly comparable to those of the house fixed-effect model. $\psi_{t,t'}$ replaces the intercept in equation

(4), $\psi_{t',t} = \alpha_{t'} - \gamma(t,t')\alpha_t$, and $Z'\delta_{t,t'}$ controls flexibly for any attributes that do not vary over time.[6] This implies the following specification:

Efficient Housing Market Model

$$\ln(p_{j,3}) = \psi_{2,3} + \gamma(2,3)\ln(p_{j,2}) - x'_{j,2}\gamma(2,3)\beta + x'_{j,3}\beta + z'_j\delta_{2,3} + \eta_{j,2,3},$$
$$\ln(p_{j,3}) = \psi_{1,3} + \gamma(2,3)\gamma(1,2)\ln(p_{j,1}) - x'_{j,1}\gamma(2,3)\gamma(1,2)\beta + x'_{j,3}\beta + z'_j\delta_{1,3} + \eta_{j,1,3},$$
$$\ln(p_{j,2}) = \psi_{1,2} + \gamma(1,2)\ln(p_{j,1}) - x'_{j,1}\gamma(1,2)\beta + x'_{j,2}\beta + z'_j\delta_{1,2} + \eta_{j,1,2},$$

where the subscripts $\{1,2,3\}$ correspond to each of the three time periods, $x'_t \equiv \{PM10, SO2, O3\}$ and z includes the housing attributes described in Table 1 and a vector of county fixed effects. Depending upon in which two of these time periods a particular house sells, one of these three equations will apply to it. The first equation applies when $t = 2$ and $t' = 3$, the second equation applies when $t = 1$ and $t' = 3$, and the third equation applies when $t = 1$ and $t' = 2$.

Given the linearity of the hedonic pricing equation, we implement a 2SLS approach to deal with the endogeneity of $x_{t'}$. In particular, we first estimate the following regression equations using all the exogenous and predetermined variables as instruments:

$$x_{j,3} = \Pi_0 Year_j + \Pi_1(3,2)x_{j,2} + \Pi_2(3,2)\ln(p_{j,2}) + \Pi_3(3,2)County_j + \upsilon_{j,2,3},$$
$$x_{j,3} = \Pi_0 Year_j + \Pi_1(3,1)x_{j,1} + \Pi_2(3,1)\ln(p_{j,1}) + \Pi_3(3,1)County_j + \upsilon_{j,1,3},$$
$$x_{j,2} = \Pi_0 Year_j + \Pi_1(2,1)x_{j,1} + \Pi_2(2,1)\ln(p_{j,1}) + \Pi_3(2,1)County_j + \upsilon_{j,1,2}.$$

where $County$ denotes a vector of county dummies and $Year$ denotes a vector of year dummies indicating t'. The first equation applies to houses that sell in periods $t = 2$ and $t' = 3$, the second equation applies to houses that sell in periods $t = 1$ and $t' = 3$, and the third equation applies to houses that sell in periods $t = 1$ and $t' = 2$. We then use the two sets of equations described above to form a GMM objective function.

Table 5 reports the results of a cross-sectional specification that considers all three pollutants simultaneously, along with specifications that consider each pollutant individually. For many pollutant-year combinations, MWTP exhibits the counterintuitive (i.e., positive) sign. Moreover, for every pollutant, results are unstable across years. These results suggest the presence of (possibly

[6] In particular, if $z'_j\phi_t$ represents the contribution of time-invariant attributes z_j to $\ln p_{j,t}$, then $z'_j\delta_{t,t'} = z'_j(\phi_{t'} - \gamma(t,t')\phi_t)$. For convenience, we label $\delta_{t,t'} = \phi_{t'} - \gamma(t,t')\phi_t$.

Table 5: Implicit Price of Pollution: Cross Sectional Estimates (N=74,892)

	Period 1			Period 2			Period 3		
	Coeff	Elast.[§]	WTP[‡]	Coeff	Elast.[§]	WTP[‡]	Coeff	Elast.[§]	WTP[‡]
A. Regressions Controlling for All Pollutants									
PM10 ($\mu g/m^3$)	0.0144	0.3252	548.9	0.0180	0.4055	684.6	-0.0337	-0.7605	-1283.8
	[0.0006]			[0.0009]			[0.0008]		
S02 (ppm)	2.4958	0.0877	95.0	8.7280	0.3066	332.3	-3.6673	-0.1288	-139.6
	[0.2462]			[0.3023]			[0.1232]		
O3 (ppm)	-2.2252	-0.2211	-84.7	-0.2005	-0.0199	-7.6	-2.2920	-0.2277	-87.3
	[0.1769]			[0.1103]			[0.1325]		
B. Separate Regression for Each Pollutant									
PM10 ($\mu g/m^3$)	0.0170	0.3836	647.6	0.0267	0.6032	1018.3	-0.0372	-0.8384	-1415.4
	[0.0006]			[0.0008]			[0.0008]		
S02 (ppm)	2.7123	0.0953	103.3	10.7170	0.3765	408.0	-4.7565	-0.1671	-181.1
	[0.2372]			[0.2861]			[0.1256]		
O3 (ppm)	-2.3586	-0.2343	-89.8	-0.4110	-0.0408	-15.6	-2.5400	-0.2524	-96.7
	[0.1703]			[0.1082]			[0.1157]		

Heteroskedasticity robust standard errors in brackets. Controls for lot size, square feet, number of rooms, number of bedrooms, number of bathrooms, year built and county fixed effects also included but not reported. [§] Elasticities calculated at means of pollutants, which are 22.55 for PM10, 0.0351 for SO2, 0.0994 for O3. [‡] Willingness to pay calculated for marginal 1 $\mu g/m^3$ change in PM10 and 1 ppb change in other pollutants, annualized at rate of 0.07 for average house price of $ 543,896.

Table 6: Implicit Price of Pollution: Fixed Effect Estimates (N=74,892)

	All Pollutants†			Single Pollutant†		
	Coeff	Elast.§	WTP‡	Coeff	Elast.§	WTP‡
PM10	0.0035	0.0790	133.3	0.0031	0.0695	117.3
($\mu g/m^3$)	[0.0003]			[0.0003]		
S02	-1.0467	-0.0368	-39.9	-0.9374	-0.0329	-35.7
(ppm)	[0.0787]			[0.0762]		
O3	-1.6736	-0.1663	-63.7	-1.7587	-0.1747	-67.0
(ppm)	[0.0630]			[0.0627]		

Heteroskedasticity robust standard errors in brackets. Controls for lot size, square feet, number of rooms, number of bedrooms, number of bathrooms, year built and county fixed effects also included but not reported. † The single pollutant regressions are run separately for each pollutant, whereas the other estimates are run with all pollutants in a single regression. § Elasticities calculated at means of pollutants, which are 22.55 for PM10, 0.0351 for SO2, 0.0994 for O3. ‡ Willingness to pay calculated for marginal 1 $\mu g/m^3$ change in PM10 and 1 ppb change in other pollutants, annualized at rate of 0.07 for average house price of $ 543,896.

Table 7: Implicit Price of Pollution: Efficient Markets (N=74,892)

	Three Pollutants Simultaneously						Single Pollutant†		
	Without Instruments*			With Instruments*			With Instruments*		
	Coeff	Elast.§	WTP‡	Coeff	Elast.§	WTP‡	Coeff	Elast.§	WTP‡
PM10	-0.0082	-0.1844	-311.2	-0.0104	-0.2345	-396.0	-0.0132	-0.2977	-502.6
($\mu g/m^3$)	[0.0003]			[0.0004]			[0.0004]		
S02	-2.3851	-0.0838	-90.8	-4.2285	-0.1484	-161.0	-5.1522	-0.1808	-196.2
(ppm)	[0.0790]			[0.1233]			[0.1176]		
O3	-2.1489	-0.2135	-81.8	-2.5737	-0.2558	-98.0	-2.8211	-0.2804	-107.4
(ppm)	[0.0687]			[0.1237]			[0.1102]		

Heteroskedasticity robust standard errors in brackets. Controls for prior sales price, lot size, square feet, number of rooms, number of bedrooms, number of bathrooms, year built and county fixed effects also included but not reported. ∗ with instruments indicates that instrument for future pollution with measures at time of purchase. † The single pollutant regressions are run separately for each pollutant, whereas the other estimates are run with all pollutants in a single regression. § Elasticities calculated at means of pollutants, which are 22.55 for PM10, 0.0351 for SO2, 0.0994 for O3. ‡ Willingness to pay calculated for marginal 1 $\mu g/m^3$ change in PM10 and 1 ppb change in other pollutants, annualized at rate of 0.07 for average house price of $ 543,896.

Table 8: Implicit Price of House Attributes (N=74,892)

	Efficient Market IV		Fixed Effect		Cross Section	
	Coeff	Elast.§	Coeff	Elast.§	Coeff	Elast.§
	Sold in period 3 and 2				Sold in period 3	
Lot size	4.52E-06*	0.0316	2.15E-07	0.0015	1.40E-05*	0.0978
	[2.51E-07]		[2.63E-07]		[4.60E-07]	
Square feet	1.15E-04*	0.1978	-5.46E-05*	-0.0939	3.52E-04*	0.6055
	[3.84E-06]		[3.78E-06]		[4.55E-06]	
No. of bedrooms	5.40E-03*	0.0107	1.70E-02*	0.0550	1.84E-03	0.0060
	[2.00E-03]		[2.25E-03]		[2.40E-03]	
No. of rooms	8.55E-04	0.0028	-6.10E-03	-0.0416	7.47E-03*	0.0510
	[9.73E-04]		[-6.10E-03]		[1.57E-03]	
No. of bathrooms	1.34E-02*	0.0915	-6.02E-03*	-0.0120	4.27E-02*	0.0850
	[2.70E-03]		[3.05E-03]		[3.15E-03]	
Year built	-2.10E-03*	-4.1307	-1.65E-03*	-3.2438	-1.80E-03*	-3.5412
	[6.82E-05]		[7.52E-05]		[8.58E-05]	
	Sold in period 3 and 1				Sold in period 2	
Lot size	6.37E-06*	0.0445	9.65E-07*	0.0067	1.24E-05*	0.0866
	[4.50E-07]		[4.09E-07]		[4.30E-07]	
Square feet	1.80E-04*	0.3096	-3.24E-05*	-0.0557	4.50E-04*	0.7741
	[5.87E-06]		[6.70E-06]		[5.80E-06]	
No. of bedrooms	3.90E-03	0.0078	1.48E-02*	0.0479	-2.86E-02*	-0.0926
	[3.20E-03]		[3.88E-03]		[3.25E-03]	
No. of rooms	2.90E-03*	0.0094	-2.43E-03	-0.0166	1.45E-02*	0.0990
	[1.10E-03]		[2.48E-03]		[2.14E-03]	
No. of bathrooms	3.26E-02*	0.2226	1.10E-02*	0.0219	3.68E-02*	0.0732
	[4.60E-03]		[4.87E-03]		[4.29E-03]	
Year built	-2.80E-03*	-5.5076	-2.61E-03*	-5.1348	-1.02E-03*	-2.0067
	[1.19E-04]		[1.28E-04]		[1.06E-04]	
	Sold in period 2 and 1				Sold in period 1	
Lot size	2.51E-06*	0.0175	1.04E-06*	0.0073	1.20E-05*	0.0838
	[4.39E-07]		[3.56E-07]		[5.40E-07]	
Square feet	8.05E-05*	0.1385	7.45E-06	0.0128	4.42E-04*	0.7603
	[6.21E-06]		[4.68E-06]		[6.96E-06]	
No. of bedrooms	-2.04E-04	-0.0004	5.20E-03	0.0168	-2.21E-02*	-0.0715
	[3.40E-03]		[2.91E-03]		[3.84E-03]	
No. of rooms	6.36E-04	0.0021	-2.27E-04	-0.0015	8.18E-03*	0.0558
	[8.77E-04]		[7.27E-04]		[2.24E-03]	
No. of bathrooms	6.70E-03	0.0457	7.89E-03	0.0157	3.44E-02*	0.0684
	[4.80E-03]		[4.21E-03]		[5.43E-03]	
Year built	-1.10E-03*	-2.1637	-1.26E-03*	-2.4789	-2.02E-04	-0.3974
	[1.21E-04]		[1.07E-04]		[1.41E-04]	

Heteroskedasticity robust standard errors in brackets. ∗ indicate statistically significantly different from 0 at 95% level. These parameter estimates are taken from the same regressions for which the pollutant coefficients are reported in Table 7 (column 4), Table 6 (column 1), and Table 5 (panel A, columns 1,4,and 7). § Elasticities calculated at means of house characteristics as reported in Table 1.

time-varying) unobservables that are correlated with the pollutants we are studying.

The fixed-effect model uses panel variation to address the impact of time-invariant unobservables. Table 6 describes these results. MWTP estimates for SO2 and O3 are stable across specifications, have the expected sign, and are small but not unreasonable in magnitude. MWTP for PM10, however, has a counterintuitive sign, suggesting the presence of some sort of unobservable that was not adequately controlled for by the house fixed effect. This is the typical sort of bias encountered in the hedonic valuation of air pollution – desirable unobservables may evolve over time in conjunction with worsening air pollution (e.g., the opening of new businesses, or other forms of economic growth). The house fixed effect is unable to control for this sort of evolving unobservable.

It is at this point in the research process where previous work has turned to some sort of quasi-random source of variation in pollution to accurately identify MWTP. There is no natural source of quasi-random variation in our Bay Area data set, so we instead turn to our model based on a limited notion of informational efficiency. The results of this model for the pollution variables are described in Table 7. That table reports results without instruments (i.e., omitting the 2SLS procedure used to deal with the endogeneity of current pollution levels), with instruments (i.e., using a non-linear 2SLS approach to deal with the endogeneity of contemporaneous pollution), and considers all pollutants both simultaneously and one at a time. Results without the 2SLS procedure are reported in order to make clear that the majority of the difference between our results and the results of the fixed-effects model do not arise because of the use of IV, but rather because of the way in which we treat the unobservable product attribute.

We first consider the results for PM10 in detail. Whereas the fixed-effect estimates of the MWTP for PM10 had a counterintuitive sign, estimates from our efficient housing market model imply a statistically significant MWTP to avoid an additional microgram of particulate matter per cubic meter ranging between \$396 and \$502. Of the pollutants that we study, particulate matter has received the most attention in the hedonics literature, and Chay and Greenstone (2005) provide a good set of comparison results. They measure the value of total suspended particulate (TSP) reductions by looking at how the median house price in a county varies with TSP concentration, assuming a national housing market and ignoring the wage gradient described in Roback (1982).[7]

[7] Prior to 1987, the EPA measured the concentration of a wide range of particulate matter of various sizes, denoted by total suspended particulates (TSP). After 1987, the EPA switched its focus to "inhalable coarse particles" with diameters between 2.5 and 10 micrometers, and "fine particles" with diameters less than 2.5 micrometers. PM10 refers to any particle with a diameter smaller than 10 micrometers. These particles, which are the focus of our analysis, are considered to have greater adverse health consequences because of their potential to travel deep into the lungs and even into the bloodstream. (http://www.epa.gov/air/particlepollution/basic.html)

Similar to our results, using cross-sectional evidence from 1970 and 1980, they find statistically weak correlations between TSP and county median house prices, and the sign of the estimated effect varies with the cross-sectional specification.[8] Similar results are obtained even when fixed effects are used to control for permanent unobserved differences between counties. This suggests a need for controls to deal with time-varying unobservables. Because Chay and Greenstone treat the US as a single housing market, they are able to observe variation in EPA attainment status for TSP across locations within a year. When the EPA declares an area to be out of attainment, local officials are required to implement strict regulations to bring pollution levels down. Using a variety of tests, Chay and Greenstone find that this creates an exogenous source of variation in TSP (i.e., one that only affects housing values through its effect of reducing TSP); TSP attainment status can therefore be used as an instrument. However, because we restrict our attention to a more conventionally defined housing market (i.e., a single metropolitan area – the San Francisco Bay Area), and since that area is treated as a single entity for the purposes of air pollution regulation, within-year variation in attainment status is not available to us for use as an instrument. We instead rely on the efficient housing market hypothesis. Using their instrumental variables strategy, Chay and Greenstone find a housing price elasticity with respect to TSP between -0.2 and -0.35. Those results are statistically significant and substantially larger than previous results in the literature.[9] Our results (in particular, when we consider PM10 in isolation, as is done for TSP by Chay and Greenstone) are quite similar; we find an elasticity of -0.30. This number falls to -0.23 when we consider all three pollutants together.

Similar biases appear to be present for O3 and SO2, although in neither case is the bias as severe as in the case of PM10. In the case of SO2, MWTP rises from \$36 to \$296 when time-varying unobservables are accounted for. In the case of O3, MWTP only rises from \$67 to \$107, suggesting that, while still an issue, time-varying unobservables may not be as serious of a concern for this pollutant.

Table 8 describes the results of all three models for non-pollution housing attributes. For the efficient housing and fixed-effects models, these estimates describe the difference in the implicit price of each attribute over time (e.g., the difference between the period 3 and period 2 coefficients on lot-size in the efficient housing model is a statistically significant $4.52x10^{-6}$). We see that the

[8]The greater statistical significance in our application is likely the result of the much larger sample size in our housing transactions data set, compared to Chay and Greenstone's data set which used counties as the unit of observation.

[9]Previous papers that did not control for time-varying unobservables (many of which ignored time-invariant unobservables as well) found smaller (often counter-intuitively signed) marginal willingnesses to pay. Smith and Huang (1995) survey the literature from 1967 to 1988 that values marginal reductions in particulate matter in the context of a meta-analysis. They find elasticities that tend to lie between -0.04 and -0.07.

implicit prices of many attributes associated with larger homes tend to rise over time under the efficient housing model, while they more often fall under the fixed-effects model (although this is by no means uniform across all attributes and many of the estimates are insignificant). The value of newer homes (i.e., year-built) falls over time in both of these specifications. The cross-sectional coefficient estimates for non-pollution housing attributes in each period can be easily interpreted as implicit prices.

5 Conclusion

Our paper demonstrates a new approach to controlling for unobserved product attributes in hedonic models. In particular, we show how a necessary condition for a market to be informationally efficient can be exploited to identify implicit prices in the context of either fixed or time-varying unobserved product attributes. We then describe an estimator that can be applied to settings where repeat sales data are available and our limited form of informational efficiency is likely to hold.

We use our estimator to recover a consumer's marginal willingness to pay for clean air in the Bay Area. Particularly appealing features of our identification strategy are that (i) it can be easily applied to data from a single, well-defined housing market, and (ii) it relies on a testable assumption (i.e., that available information should not predict economically significant excess returns). We find evidence that this assumption is valid for the housing market in the Bay Area.

We estimate the implicit price of three of the EPA's criteria air pollutants (PM10, SO2, and O3). In contrast to fixed effects methods (which just control for time-invariant unobservables) or cross sectional methods (which ignore correlated unobservable attributes altogether), our estimates of the implicit price indicate that consumers value pollution reductions, and that their MWTP to avoid pollution is significantly larger in magnitude than that found by other models. Particularly in the case of PM10, it appears that failing to control for unobserved attributes at all or only controlling for time-invariant unobserved attributes leads to the wrong sign on the estimate of the potential benefits of a pollution reduction policy. Our estimates suggest that SO2 is also prone to a large bias from ignoring time-varying unobservables. Time varying unobservables associated with ground-level ozone, while important, do not appear to lead to as large of a bias.

To be clear, while our approach works well in this context, we are not claiming that it will be superior to quasi-random approaches in all applications. The identifying assumptions in our approach and quasi-random approaches are not nested, and the plausibility of either set of assumptions depends on the particular application and data one is using. Our approach may be preferable when a legitimate source of quasi-randomness cannot be found, or when one is available but it generates insufficient exogenous variation in the variable of interest. On the other hand, quasi-randomness will be preferable if there is reason to suspect that the housing market being studied was failing to function efficiently. In empirical work, we advise applied researchers to test the sensitivity of results to alternative identifying assumptions when they are available.

References

Ai, C. and X Chen (2003), "Efficient Estimation of Models With Conditional Moment Restrictions Containing Unknown Functions," *Econometrica*, 71(6), 1795-1843.

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

Black, SE (1999), "Do Better Schools Matter? Parental Valuation of Elementary Education," *The Quarterly Journal of Economics*, 114(2): 577-599.

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

Chay, KY and M Greenstone (2005), "Does Air Quality Matter? Evidence from the Housing Market," *Journal of Political Economy*, 113(2), 376-424.

D'Ippoliti D, F Forastiere, C Ancona, N Agabity, D Fusco, P Michelozzi, and CA Perucci (2003), "Air Pollution and Myocardial Infarction in Rome: A Case-Crossover Analysis," *Epidemiology*, 14: 528-535.

Dockery DW, CA Pope, X Xu, JD Spengler, JH Ware, ME Fay, BG Ferris, and FE Speizer (1993), "An Association Between Air Pollution and Mortality in Six U.S. Cities," *New England Journal of Medicine*, 329: 1753-1759.

Dougherty, A. and R Van Order (1982), "Inflation, Housing Costs, and the Consumer Price Index," *The American Economic Review*, 72(1): 154-164.

Ekeland, I, JJ Heckman, and L Nesheim (2004), "Identification and Estimation of Hedonic Models," *Journal of Political Economy*, 112(S1): 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.

Ghio AJ, C Kim, and RB Devlin (2000), "Concentrated Ambient Air Particles Induce Mild Pulmonary Inflammation in Healthy Human Volunteers," *American Journal of Respiratory and Critical Care Medicine*,162(3 Pt.1): 981-988.

Heckman, J, RL Matzkin, and L Nesheim. "Simulation and Estimation of Hedonic Models," Center for Economic Studies & Ifo Institute for Economic Research (CESifo), Working Paper No. 1014, August 2003.

Hong YC, JT Lee, H Kim, EH Ha, J Schwartz, and DC Christiani (2002), "Effects of Air Pollutants on Acute Stroke Mortality," *Environmental Health Perspectives*, 110: 187-191.

Lin M, Y Chen, RT Burnett, PJ Villeneuve, and D Kerwski (2002), "The Influence of Ambient Coarse Particulate Matter on Asthma Hospitalization in Children: Case-Crossover and Time-Series Analyses," *Environmental Health Perspectives*, 110: 575-581.

Murphy, K and R Topel (1985), "Estimation and Inference in Two-Step Econometric Models," *The Journal of Business & Economic Statistics*, 3-4, 370-379.

Norris G, SN YoungPong, JQ Koenig, TV Larson, L Sheppard, and JW Stout (1999), "An Association Between Fine Particles and Asthma Emergency Department Visits for Children in Seattle," *Environmental Health Perspectives*, 107: 489-493.

Pope CA, RT Burnett, MJ Thun, EE Calle, D Krewski, K Ito, and GD Thurston (2002), "Lung Cancer, Cardiopulmonary Mortality, and Long-Term Exposure to Fine Particulate Air Pollution," *JAMA*, 287:1132-1141.

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

Samet JM, DM DeMarini, and HV Malling (2004), "Do Airborne Particulates Induce Heritable Mutations?" *Science*, 304(5673):971-972.

Slaughter JC, T Lumley, L Sheppard, JQ Koenig, and GG Shapiro (2003), "Effects of Ambient Air Pollution on Symptom Severity and Medication Use in Children with Asthma," *Annals of Allergy, Asthma, and Immunology*, 91: 346-353.

Smith, VK and JC Huang (1995), "Can Markets Value Air Quality? A Meta-Analysis of Hedonic Property Value Models," *Journal of Political Economy*, 103(1), 209-227.

Tolbert PE, JA Mulholland, DD MacIntosh, F Xu, D Daniels, OJ Devine, BP Carlin, M Klein, J Dorley, AJ Butler, DF Nordenberg, H Franklin, PB Ryan, and MC White (2000), "Air Quality and Pediatric Emergency Room Visits for Asthma in Atlanta, Georgia," *American Journal of Epidemiology*, 151: 798-810.

Tsai SS, WB Goggins, HF Chiu, and CY Yang (2003), "Evidence for an Association Between Air Pollution and Daily Stroke Admissions in Kaohsiung, Taiwan," *Stroke*, 34(11): 2612-2616.