Incorporating Dynamic Behavior into the Hedonic Model*

Kelly C. Bishop
Arizona State University

Alvin D. Murphy
Arizona State University

Fall 2019

Preliminary. Do not circulate. Comments welcome.

Abstract

The property value hedonic model, which has been used extensively in the environmental and urban literature, has long been considered the "workhorse" model of amenity valuation. Derived using the household's first-order conditions associated with choosing where to live, this model is both intuitive and straightforward to estimate. However, a recognized drawback of the existing literature is that the model assumes a static framework (i.e., it assumes that there are no costs associated with moving and that agents do not look to the future when choosing where to live today). This is clearly an unrealistic assumption; real estate fees alone usually stand at six percent of a property's sales price and households typically have strong emotional ties to both their home and neighborhood. Facing such high costs, households will behave dynamically when choosing where to live today (and when choosing how much of each housing amenity to consume). In this paper, we develop an approach that incorporates dynamic behavior into the traditional Rosen-style hedonic model of demand. In an application, we use this dynamic model to estimate the willingness to pay to avoid violent crime and air pollution and find evidence of significant bias in the static model.

Key Words: Hedonic Demand, Willingness to Pay, Valuation, Housing, Dynamics
JEL Classification Numbers: Q50, Q51, R21, R23

*We would like to thank Dionissi Aliprantis, Soren Anderson, Peter Arcidiacono, Peter Blair, Leah Brooks, Nick Kuminoff, Dennis Epple, Bob Miller, Aviv Nevo, Christopher Palmer, Monika Piazzesi, Chris Taber, Chris Timmins, Matt Wiswall, and participants at University of Toronto, University of Western Ontario, Cornell University, Federal Reserve Board of Governors, George Washington Univeristy, University of Wisconsin–Madison, University of Houston, University of Nevada – Las Vegas, NBER Summer Institute, Heartland Workshop at Illinois, NARSC Annual Meetings, AERE Winter Meetings, Econometric Society Summer Meetings, AREUEA Winter Meetings, HULM conference, and WEAI Annual Meeting for valuable comments on this paper. All remaining errors are our own. Email: kelly.bishop@asu.edu, alvin.murphy@asu.edu

1 Introduction

The property value hedonic model, which has been used extensively in the environmental and urban literature, has long been considered the "workhorse" model of amenity valuation.[1] Based on the original framework laid out in Rosen's seminal paper, "Hedonic Prices and Implicit Markets" (1974), the model is derived using the household's first-order conditions associated with choosing where to live and is both intuitive and straightforward to estimate. However, a recognized drawback of the existing hedonic literature is that the models assume a static framework (i.e., they assume that there are no moving costs, no wealth effects, and that agents are myopic). This is clearly an unrealistic assumption, as changing residence is known to be a costly undertaking; real estate fees alone usually stand at six percent of a property's sales price and households typically have strong emotional ties to both their home and neighborhood. In the face of such costs, households will behave dynamically when choosing where to live today (and how much amenities to consume). In this paper, we develop an approach that incorporates dynamic behavior into the traditional Rosen-style hedonic model of demand.[2]

Our model expands the existing hedonic framework by allowing households to be forward-looking with respect to the amenities of interest and these amenities' implicit prices. Specifically, we allow for the fact that a change in the consumption level of an amenity affects not only current utility (through both increased price and increased amenity consumption), but also the stream of future utilities. Thus, we re-define the familiar optimality condition

[1] See, for example, applications focusing on school quality (Black (1999), Downes and Zabel (2002), Gibbons and Machin (2003)), climate (Albouy, Graf, Kellogg, and Wolff (2016)), safety (Gayer, Hamilton, and Viscusi (2000), Davis (2004), Greenstone and Gallagher (2008)), and environmental quality (Palmquist (1982), Chay and Greenstone (2005), Bento, Freedman, and Lang (2015)

[2] For recent papers that estimate dynamic discrete choice models of location choice, see Kennan and Walker (2011), Bishop (2012), Bayer, McMillan, Murphy, and Timmins (2016), Caetano (2016), Davis, Gregeory, Hartley, and Tan (2017), and Mastromonaco (2015). Bishop and Murphy (2019) develops a simple, forward-looking hedonic model that does not allow for re-optimization and presents analytical results which relate the direction of bias to the time-series properties of the amenity of interest. Recent papers have estimated dynamic models of housing supply, including Paciorek (2013), Murphy (2018), and Mangum (2017).

to require that the marginal increase in current price is equal to the marginal benefit of an increase in current amenity consumption *plus* the associated change in future utility flow.

Our dynamic estimator, which combines insights from the recent dynamic discrete-choice literature described (Hotz and Miller (1993) and Arcidiacono and Miller (2011)) with those in the recent hedonic literature (Ekeland, Heckman, and Nesheim (2004) and Bishop and Timmins (2017)), allows for both transaction costs and forward-looking behavior of the households and adds relatively little to the computational burden; we require only the additional first-stage estimation of the change in future utility flows associated with an amenity change today. We show that once this additional first stage is estimated, the utility function parameters may be recovered using existing methods derived for the static framework. As our estimator has a low computational burden, we are able to additionally control for both unobserved house/neighborhood attributes, as well as unobserved individual preference heterogeneity within our dynamic framework.

In an application, we estimate the willingness to pay to avoid crime and pollution in the Bay Area. We do this using an unusually rich set of data describing both housing transactions and time-varying household and neighborhood attributes, including house-level measures of crime and air pollution. Our data is constructed from a number of sources. The first component is data on all housing transaction in the Bay Area from 1990-2008, which provides rich information on housing characteristics and sales prices. It also provides a unique house identifier and buyer and seller names, which can be used to partially track households through time. The second component is demographic information on buyers, which is obtained from mortgage applications. The third component is geo-coded amenity data on violent crime and pollution.

Using these data, we estimate that the mean household is willing to pay \$382 per year to reduce violent crime by 10% and willing to pay \$379 per year to reduce pollution (as measured by PM10) by 10%. Furthermore, we find that ignoring dynamics (i.e., by running the standard myopic model) imposes substantial downward biases on the estimated willingness to pay to avoid violent crime and the estimated willingness to pay to avoid air pollution.

This paper proceeds as follows. In Section 2, we introduce our dynamic model and show how it may be "broken up" into a two-part, discrete/continuous-choice representation of the model. In Section 3, we develop our estimation strategy, which takes advantage of this two-part decision framework. In Section 4 and 5, we present our data and the results from our application of this model using data describing the Bay Area of California. Finally, Section 6 concludes.

2 A Dynamic Model of Hedonic Demand

In the traditional hedonic framework, households make a single amenity-consumption decision to maximize current utility. Our dynamic framework models households as making a sequence of amenity-consumption decisions to maximize the discounted sum of expected per-period utilities.

Households (denoted $i \in \{1, \ldots, N\}$) have heterogenous preferences over a vector of housing-related amenities, $x_{i,t}$ (which may be fixed, such as lot size, or time-varying, such as local crime rates). We allow each household's preferences and moving costs to differ based on a vector of observable attributes, $z_{i,t}$ (which may be fixed, such as race, or time-varying, such as income). We also allow preferences to differ across households by vector of unobserved household- and amenity-specific preference shocks, $\eta_{i,t}$, and allow moving cost to vary by an unobserved moving cost shock, $\epsilon_{i,t}$.[3]

[3] The treatment and timing of $\eta_{i,t}$ with respect to the estimation of this model is discussed in detail in Section 3.3.

Households begin each period t with an endowment vector of x, which is determined by its current residence and is denoted $x^e_{i,t}$. In each period, households then choose how much of each amenity in the vector $x_{i,t}$ to consume. This choice is continuous over the support of x. If the household decides to consume an amenity level different from that offered by its current residence, the household must incur a moving cost and move to a house which offers its optimal level of amenities. When this moving cost is sufficiently high (relative to the potential gains from moving), the household will choose to stay in its current residence and consume their endowed quantity of amenities.

It is convenient to specify households as facing a two-part discrete-continuous choice in each period: the household first decides whether or not to move and, conditional upon moving, how much of each amenity to consume.[4] This also facilitates the two-step estimation strategy described in the next section of the paper. The first decision of whether or not to move (and incur the associated moving cost), is discrete and is denoted $d^m_{i,t} \in \{0, 1\}$. The second decision, of how much x to consume, is continuous over the support of x and is denoted $d^x_{i,t}$. The second, continuous-choice decision is made only if $d^m_{i,t} = 1$; if the household chooses not to move in a given period ($d^m_{i,t} = 0$), it makes no further decision and consumes its endowment level of amenities (i.e., $d^x_{i,t} = x^e_{i,t}$).

We write the household's problem as choosing both $d^m_{i,t}$ and $d^x_{i,t}$ to maximize the expected discounted sum of per-period flow utilities, which is given by:

$$E\Big[\sum_{\tau=t}^{T} \beta^{\tau-t}\big(u^f_{\mathrm{m}}(s_{i,\tau}, x_{i,\tau}, \epsilon_{i,\tau})\big)|s_{i,t}, d^m_{i,t} = \mathrm{m}, d^x_{i,t} = x_{i,t}\Big] \tag{1}$$

[4]See Dubin and McFadden (1984) and Hanemann (1984) for seminal papers that estimate discrete-continuous demand models.

where:

$$
\begin{aligned}
u_0^f(s_{i,t}) &= u(x_{i,t}^e, z_{i,t}, \eta_{i,t}; \alpha) - r(x_{i,t}^e; \gamma^r) \\
u_1^f(s_{i,t}, x_{i,t}, \epsilon_{i,t}) &= u(x_{i,t}, z_{i,t}, \eta_{i,t}; \alpha) - r(x_{i,t}; \gamma^r) - MC(z_{i,t}; \gamma^M) - \epsilon_{i,t}
\end{aligned} \tag{2}
$$

and where the state vector, $s_{i,t}$, is comprised of all state variables (except $\epsilon_{i,t}$) that affect the household's decisions, $d_{i,t}^m$ and $d_{i,t}^x$. Thus, $s_t = [x_{i,t}^e, z_{i,t}, \eta_{i,t}, \Omega_{i,t}]$, where $\Omega_{i,t}$ is household i's information set at time t. We assume that the shock to moving costs, $\epsilon_{i,t}$, has no predictive power for $s_{i,t+1}$, conditional on $s_{i,t}$ and $d_{i,t}$[5] and that it is distributed *i.i.d.* over both households and time according to the Logistic distribution with scale parameter σ_ϵ.

The choice-specific flow utilities are defined for each possible realization of the discrete choice, $d_{i,t}^m = \mathrm{m}$, where $\mathrm{m} \in \{0, 1\}$ and are comprised of three components. The first component of utility, $u(x_{i,t}, z_{i,t}, \eta_{i,t}; \alpha)$, is known up to the parameter vector α and is a function of consumed amenities, $x_{i,t}$, and not a function of endowed amenities, $x_{i,t}^e$ (although if the household chooses to not move, $x_{i,t} = x_{i,t}^e$). The second component is the implicit rental cost of x, $r(x_{i,t}; \gamma^r)$, which is known up to the parameter vector, γ^r. This rental price component, (entering with a negative sign) shows that a higher consumption of $x_{i,t}$ reduces consumption of other goods. The third component is the moving cost, $MC(z_{i,t}; \gamma^M) + \epsilon_{i,t}$, which is paid only if the household moves and is known up to the parameter vector, γ^M. This moving cost component reflects the fact that the household may choose to consume any quantity of $x_{i,t}$, but will need to pay the moving cost if it chooses a level of $x_{i,t}$ other than $x_{i,t}^e$. In the standard Rosen framework, there is no endowment of amenities and no cost associated with reoptimization.

If the household chooses to move, it knows that it will optimally make the decisions $d_{i,t}^m$

[5]This is analogous to the familiar Conditional Independence assumption made in Rust (1987) and allows us to write the transition densities as $q = q_s(s_{i,t+1}|s_{i,t}, d_{i,t})q_\epsilon(\epsilon_{i,t+1})$.

and $d^x_{i,t}$ in all future periods. This allows us to define the direct choice-specific value function associated with moving, $v_1(s_{i,t}, x_{i,t})$. This function specifies the lifetime utility that a household will receive from choosing to move in the current period (given any subsequent choice of $x_{i,t}$). Therefore, if the household chooses to move, $x_{i,t}$ is chosen to maximize $v_1(s_{i,t}, x_{i,t})$:

$$x^*_{i,t}(s_{i,t}) = argmax_x \; v_1(s_{i,t}, x_{i,t})$$

The household also knows that it will behave optimally in the future if it chooses to not move in the current period. However, when choosing to not move, there is no subsequent choice of $x_{i,t}$ and the household consumes its endowment level of amenities, $x^e_{i,t}$.

Given the respective amenity-choice outcomes associated with each of the discrete-choice alternatives, we may now define the indirect choice-specific value functions for not moving and moving, $v_0(s_{i,t})$ and $v_1(s_{i,t})$, respectively. The function $v_0(s_{i,t})$ recognizes that the endowment vector, $x^e_{i,t}$, is an element of $s_{i,t}$. The function $v_1(s_{i,t})$ is defined assuming that $x_{i,t}$ is chosen optimally, i.e., $v_1(s_{i,t}) = v_1(s_{i,t}, x^*_{i,t}(s_{i,t}))$, and is no longer a function of $x_{i,t}$ itself (but rather a function of the variables that determine the optimal $x^*_{i,t}(s_{i,t})$ only). These indirect choice-specific value functions are given by:

$$\begin{aligned} v_0(s_{i,t}) &= u^f_0(s_{i,t}) + \beta E[max\{v_0(s_{i,t+1}), v_1(s_{i,t+1}) - \epsilon_{i,t+1}\}|s_{i,t}, d^m_{i,t} = 0, d^x_{i,t} = x^e_{i,t}] \\ v_1(s_{i,t}) &= \bar{u}^f_1(s_{i,t}) + \beta E[max\{v_0(s_{i,t+1}), v_1(s_{i,t+1}) - \epsilon_{i,t+1}\}|s_{i,t}, d^m_{i,t} = 1, d^x_{i,t} = x^*_{i,t}(s_{i,t})] \end{aligned}$$

where, following convention, the value functions are defined excluding the current period's shock to moving costs. For notational convenience, we let $\bar{u}^f_1(s_{i,t})$ denote the flow utility associated with moving absent the shock to moving costs, $\epsilon_{i,t}$. We do not need to denote an equivalent term for the flow utility associated with not moving, as moving costs (including the idiosyncratic shock) are only paid if the household, in fact, moves (i.e., in this case, an analogously defined $\bar{u}_0^f(s_{i,t})$ would exactly equal $u_0^f(s_{i,t})$). Note that the function $\bar{u}_1^f(s_{i,t})$ is an indirect flow utility as it is defined assuming that $x_{i,t}$ is chosen optimally, *i.e.*, $\bar{u}_1^f(s_{i,t}) = \bar{u}_1^f(s_{i,t}, x_{i,t}^*(s_{i,t}))$.

The household will choose to move if $v_1(s_{i,t}) - \epsilon_{i,t} > v_0(s_{i,t})$. If the household chooses to move ($d_{i,t}^m = 1$), $x_{i,t}$ is chosen to maximize the associated direct value function, $v_1(s_{i,t}, x_{i,t})$, which we may now write as:

$$v_1(s_{i,t}, x_{i,t}) = \bar{u}_1^f(s_{i,t}, x_{i,t}) + \beta E[max\{v_0(s_{i,t+1}), v_1(s_{i,t+1}) - \epsilon_{i,t+1}\}|s_{i,t}, d_{i,t}^m = 1, d_{i,t}^x = x_{i,t}] \quad (3)$$

3 Estimation of the Model

Estimation of the traditional, static model begins with the separate first-stage estimation of the implicit price of the amenity of interest, i.e., the estimation of the gradient of the housing price function relating amenity levels to housing prices/rents. In the second stage, the parameters of the utility function are recovered (treating as known the first-stage estimates of the hedonic price gradient).[6] Our estimation framework retains this simple intuition of the Rosen framework and adapts it for a dynamic context by adding an extra estimation stage to capture the impact of this period's amenity choice of future utility streams.

The estimation of dynamic models can come with a substantial computational burden.

[6]Because of the well-documented difficulties associated with estimating Rosen's second stage (Brown and Rosen (1982), Mendelsohn (1985), Bartik (1987), and Epple (1987)), much of the previous literature has forewent estimation past the first-stage and focused only on estimating the local effects of a policy change. However, more recent papers, such as Ekeland, Heckman, and Nesheim (2004), Bajari and Benkard (2005), Heckman, Matzkin, and Nesheim (2010), and Bishop and Timmins (2017), show how to estimate the willingness-to-pay function (*i.e.*, allow willingness-to-pay to vary with the amenity of interest) while avoiding the identification and endogeneity issues laid out in earlier papers. Nesheim (2015) and Chernozhukov, Galichon, Henry, and Pass (2017) discuss the identification of multi-amenity hedonic models. Kuminoff and Pope (2014) discusses the conditions of the hedonic model under which changes in the equilibrium price function identify willingness-to-pay parameters.

The standard, well-known computational difficulty lies in the fact that $v_0(s_{i,t})$ and $v_1(s_{i,t})$ are defined recursively. In our case, we have an additional recursive structure to contend with, as the choice-specific value function associated with moving is also a function of the optimal amenity vector, $x^*_{i,t}(s_{i,t})$, while at the same time $x^*_{i,t}(s_{i,t})$ is itself the solution to a problem involving this value function. This additional complication (a doubly-recursive structure) means that our model would be computationally prohibitive to estimate using the full-solution method described by Rust (1987).

To estimate the model, we combine insights from the recent dynamic discrete-choice literature described (Hotz and Miller (1993) and Arcidiacono and Miller (2011)) with those in the recent hedonic literature (Ekeland, Heckman, and Nesheim (2004) and Bishop and Timmins (2017)), while still retaining all of the intuition associated with the classic Rosen framework. Employing a two-step estimation approach, we show that the estimation of our dynamic model is reduced to a familiar, and computationally feasible, environment. In the first stage, we recover estimates of the price function parameters (i.e., we estimate the first stage described by Rosen) and reduced-form estimates of the future value associated with this period's choice. In the second stage, we are able to treat these first-stage estimates as "data" and recover the remaining structural parameters using existing hedonic methods.[7] As our estimation approach has a very low computational burden, we are able to additionally control for both unobserved house/neighborhood attributes, as well as very rich unobserved individual preference heterogeneity.[8]

[7]See Ekeland, Heckman, and Nesheim (2004), Heckman, Matzkin, and Nesheim (2005), Heckman, Matzkin, and Nesheim (2010), and Bishop and Timmins (2017).

[8]For an application of a simplified version of this framework with a single amenity, no wealth effects, and no unobserved heterogeneity, see Bishop and Murphy (2011).

3.1 Two-Step Simplification

We now present the details of our two-stage estimation approach to the two-choice, discrete-continuous framework. It is important to note that the model itself is not simplified; the simplification is regarding the estimation of the model.

Given the Logit assumption, we may rewrite the indirect choice-specific value functions in the familiar form:

$$
\begin{aligned}
v_0(s_t) &= u_0^f(s_{i,t}) + \beta\sigma_\epsilon E[log\Big(e^{v_0(s_{t+1})/\sigma_\epsilon} + e^{v_1(s_{t+1})/\sigma_\epsilon}\Big)|s_t, d_t^m = 0, d_t^x = x_t^e] \\
v_1(s_t) &= \bar{u}_1^f(s_{i,t}) + \beta\sigma_\epsilon E[log\Big(e^{v_0(s_{t+1})/\sigma_\epsilon} + e^{v_1(s_{t+1})/\sigma_\epsilon}\Big)|s_t, d_t^m = 1, d_t^x = x_t^*(s_{i,t})]
\end{aligned}
$$

From our model, the conditional probability of a household choosing to move is given by:

$$
P_1(s_t) = \frac{e^{v_1(s_t)/\sigma_\epsilon}}{e^{v_0(s_t)/\sigma_\epsilon} + e^{v_1(s_t)/\sigma_\epsilon}} \tag{4}
$$

where, to simplify the notation (here and for the remainder of the Section), we suppress the i subscripts.

Taking logs of Equation (4) and rearranging terms yields: $log(e^{v_0(s_t)/\sigma_\epsilon} + e^{v_1(s_t)/\sigma_\epsilon}) = v_1(s_t)/\sigma_\epsilon - log(P_1(s_t))$, which allows us to rewrite the indirect choice-specific value functions as:

$$
\begin{aligned}
v_0(s_t) &= u_0^f(s_t) + \beta\sigma_\epsilon E[-log(P_1(s_{t+1})) + v_1(s_{t+1})/\sigma_\epsilon|s_t, d_t^m = 0, d_t^x = x_t^e] \\
v_1(s_t) &= \bar{u}_1^f(s_t) + \beta\sigma_\epsilon E[-log(P_1(s_{t+1})) + v_1(s_{t+1})/\sigma_\epsilon|s_t, d_t^m = 1, d_t^x = x_t^*]
\end{aligned}
$$

Each choice-specific value function is now written as a function of the associated flow utility, next period's probability of moving, and next period's value of moving, all conditional on this period's decisions.

It is helpful to formalize the mechanism through which this period's choice determines next period's state variables directly and, therefore, determines future utility indirectly. To that end, we make the following assumption:

Assumption 1: This period's choice of x_t only affects next period's endowment, x^e_{t+1}, and doesn't affect the transition of any other state variables in s_{t+1}, i.e., $q(s'_{-x^e_{t+1}}|s'_{-x^e_t}, d^x_t = x) = q(s'_{-x^e_{t+1}}|s'_{-x^e_t})$.

Given Assumption 1, it is straightforward to show the following two results which are used to simplify the estimation of the model:

$$E[v_1(s_{t+1})|s_t, d^m_t = 1, d^x_t = x^*_t(s_t)] \quad = \quad E[v_1(s_{t+1})|s_t, d^m_t = 0, d^x_t = x^e_t] \tag{5}$$

and

$$\frac{\partial E[v_1(s_{t+1})|s_t, d^m_t = 1, d^x_t = x_t]}{\partial x_t} \quad = \quad 0 \tag{6}$$

Equation (5) states that the expected value associated with moving next period is independent of this period's discrete (move) choice, d^m_t. Equation (6) states that, conditional on moving this period, the expected value associated with moving next period is independent of this period's continuous (amenity) choice, d^x_t. Intuitively, what Assumption 1 and equations (5) and (6) say is that mobility or amenity consumption choices made in this period only affect the amenity endowment in the next period. However, if the agent will move and reoptimize consumption in the next period, their endowment doesn't matter. Therefore, this period's decisions do not determine the value of moving in the next period. They do, however, affect the value of staying in the next period.

3.2 Formation of the Likelihood

In this section, we discuss the formation of the likelihood of observing individual household choices. The two components of the joint likelihood are the likelihood contribution of the discrete choice and the likelihood contribution of the continuous choice. We outline each of these components separately and discuss under what conditions estimation can be based on only one of the two decisions.[9]

3.2.1 Continuous Choice

The likelihood contribution for the continuous choice is straightforward to form as we need only to consider the case when $d_t^m = 1$, as the continuous choice is only made when the household chooses the discrete decision of a move. If a household chooses to move, the optimal choice of x_t is given by:

$$x_t^*(s_t) = argmax_x \; v_1(s_t, x_t) - \epsilon_t = argmax_x \Big(u(x_t, z_t, \eta_t; \alpha) - r(x_t, z_t; \gamma^r) - MC(z_t; \gamma^M) - \epsilon_t - \beta \sigma_\epsilon E[log(P_1(s_{t+1}))|s_t, d_t^m = 1, d_t^x = x_t] + \beta E[v_1(s_{t+1})|s_t, d_t^m = 1, d_t^x = x_t] \Big) \quad (7)$$

Using Equation (6) (and the fact that moving costs are not a function of x_t), the optimal choice of x_t may be simplified to:

$$x_t^*(s_t) = argmax_x \Big(u(x_t, z_t, \eta_t; \alpha) - r(x_t, z_t; \gamma^r) - \beta \sigma_\epsilon E[log(P_1(s_{t+1}))|s_t, d_t^m = 1, d_t^x = x_t] \Big) \quad (8)$$

For notational convenience, we let $FV(s_t, x_t)$ describe the future-value component as-

[9] The derivation of the likelihood for the continuous choice builds upon the static hedonic likelihood developed in Heckman, Matzkin, and Nesheim (2005) and the derivation of the likelihood for the discrete choice builds upon the approach outlined in Arcidiacono and Miller (2011). One is not constrained to a likelihood-based approach. Minimum Distance or GMM estimators may also be used.

sociated with moving this period: $FV(s_t, x_t) = -E[log(P_1(s_{t+1}))|s_t, d_t^m = 1, d_t^x = x_t]$. This allows us to simplify Equation (8) to:

$$x_t^*(s_t) = argmax_x \Big(u(x_t, z_t, \eta_t; \alpha) - r(x_t, z_t; \gamma^r) + \beta\sigma_\epsilon FV(s_t, x_t) \Big) \tag{9}$$

with the first-order condition for finding the optimal choice of x_t given by:

$$u'(x_t, z_t, \eta_t; \alpha) - r'(x_t, z_t; \gamma^r) + \beta\sigma_\epsilon FV'(s_t, x_t) = 0 \tag{10}$$

Equation (10), forms the basis of the likelihood contribution of the continuous choice. For most specifications, it will be impossible to derive a closed form solution for the density of x conditional on the covariates, s_t. However, if we assume that the marginal utility function $u'(x_t, z_t, \eta_t; \alpha)$ is separable in η_t, then we can solve for a closed form solution for η_t. Once we know η_t (as a function of data and given parameter values) it is straightforward to use a change of variables to form the likelihood. Although this model is dynamic, if $FV'(s_t, x_t)$ is known, the approach to forming the likelihood outlined in Heckman, Matzkin, and Nesheim (2005) will apply.[10]

This suggests a natural two-stage estimation approach. First, we estimate the transition probabilities of the time-varying amenities and the choice probability function.[11,12] Using those estimates we can recover the "known" term $\widehat{FV}'(s_t, x_t)$, and use the change of variables (from x_t to η_t) form the likelihood. A specific example of how to construct the likelihood function is provided in the empirical specification in Section 5.1.

[10] Heckman, Matzkin, and Nesheim (2005) show how to form the likelihood in a non-parametric setting. In this case, we parameterize both the density and the utility function, however, the basic change-of-variables method of writing the likelihood still applies.

[11] For a detailed discussion of transitioning neighborhood amenities, see Lee and Lin (2018).

[12] It is straightforward to allow the price function to vary by time. This simply means that the variables predicting future γ^r are included in s.

$\widehat{FV}'(s_t, x_t)$ is given by:

$$\widehat{FV}'(s_t, x_t) = -\frac{\partial E[log(\widehat{P}_1(s_{t+1}))|s_t, d_t^m = 1, d_t^x = x_t]}{\partial x_t}$$

As $\widehat{FV}'(s_t, x_t)$ is a function of x_t, it will be correlated with the error, η_t, for reasons discussed in Epple (1987) and Bartik (1987). That is, as $x_t^*(s_t)$ is itself a function of η_t, $FV'(s_t, x_t)$ will be correlated with η_t. However, this is not a threat to identification as all we require is that η_t is independent of the state variables s_t.[13]

An analysis of Equation (10) is helpful to understand how ignoring dynamics will generate biased estimates of the marginal willingness-to-pay function. As we have normalized the coefficient on rental price to one, each household's marginal willingness to pay is given by $u'(x_t, z_t, \eta_t; \alpha)$. The static estimation will recover this as $r'(x_t, z_t; \gamma^r)$, whereas the dynamic model will recover this as $r'(x_t, z_t; \gamma^r) - \beta\sigma_\epsilon FV'(s_t, x_t)$, with the bias given by $\beta\sigma_\epsilon FV'(s_t, x_t)$.

3.2.2 Discrete Choice

As is standard in discrete-choice models, only *differences* in utility matter when estimating the dynamic discrete-choice component. Using Equation (5), the difference in utilities is simply given by:

$$\begin{aligned} &v_1(s_t) - v_0(s_t) = \bar{u}_1^f(s_t) - u_0^f(s_t) + \\ &\beta\sigma_\epsilon\big(E[-log(P_1(s_{t+1}))|s_t, d_t^m = 1, d_t^x = x_t^*(s_t)] - E[-log(P_1(s_{t+1}))|s_t, d_t^m = 0, d_t^x = x_t^e]\big) \end{aligned} \quad (11)$$

[13]If marginal willingness to pay is a function of x_t, i.e., if $u'(x_t, z_t, \eta_t; \alpha)$ is a function of x_t, then the same techniques can be used to control for the fact that x_t will be correlated with η_t in the marginal willingness-to-pay function. One could alternatively estimate (10) directly and "correct" for the endogeneity of $FV'(s_t, x_t)$ (and $u'(x_t, z_t, \eta_t; \alpha)$) by using market dummies as instruments for $FV'(s_t, x_t)$, as suggested in Bartik (1987).

Thus, once a solution for the flow utility functions, $u_0^f(s_t)$ and $\bar{u}_1^f(s_t)$, has been found,[14] forming the likelihood for the discrete choice is straightforward: first, estimate transition probabilities of each time-varying variable and the associated choice probabilities and, second, use Equations (4) and (11) to form the likelihood. As with the continuous-decision likelihood, a specific example of how to construct the likelihood function is provided in the empirical specification in Section 5.1.

3.3 Unobserved Heterogeneity

We now discuss the role of unobserved preference heterogeneity in the model's estimation, i.e., we discuss the estimation routine in light of the idiosyncratic shock to preferences, η_t. If η_t is a component of the state vector, s_t, then the first-stage estimation of the conditional choice probabilities, $P_1(s_t)$, should be conditioned upon the unobserved η_t. In the section below, we present two alternative solution methods. In both cases we assume that η_t can be written as the sum two components.

$$\eta_t = \eta_t^{pre} + \eta_t^{post} \tag{12}$$

The first component, η_t^{pre}, is observed before both the move and quantity of amenity consumptions decisions are made. The second component, η_t^{post}, is an i.i.d. shock that is observed after the agent decides to move but before deciding how much of the amenity to consume. The two main potential solutions methods reflect different assumptions that can be made about the first component, η_t^{pre}.

[14] To solve for $\bar{u}_1^f(s_t)$ one would need to know the value of $x_t^*(s_t)$ which can be found using a numerical search.

3.3.1 Simplified Model Timing

In this case, we simply assume that $\eta_t^{pre} = 0$, thereby assuming that no amenity-preference shock is observed (to the household) at the time of the household's discrete-choice decision of whether or not to move, but that the amenity-preference shock, η_t^{post}, is observed at the time of the household's continuous-choice decision of how much x_t to consume.[15] The conditional choice probability, $P_1(s_t)$, is then no longer a function of η_t and may be estimated in a completely separate first stage.

3.3.2 Discrete Heterogeneity

In this case, we specify that η_t^{pre} is discrete and follows a finite mixture distribution. As η_t^{pre} is fully observable to the household at all times, the first-stage estimation of the conditional choice probabilities, $P_1(s_t)$, needs to be conditioned upon η_t^{pre} (but not the full value of η_t). As η_t^{pre} follows a finite mixture distribution, then techniques similar to those developed in Arcidiacono and Miller (2011) may be applied. Note that in that paper, there is one unobserved variable; in this case, we are treating each element of η_t^{pre} as an unobserved state and will therefore have as many unobserved variables as number of amenities. We may allow the discrete component of η_t to be fixed through time for each household. In cases where we allow it to be time-varying, households would form expectations over future values of $\bar{\eta}_t$ in the same manner as they form expectations over other time-varying variables in the model.[16]

[15] The period t shock to moving costs, ϵ_t, is always observed at time t.

[16] Another possibility may be to assume η_t^{pre} is continuous and that $\eta_t^{post} = 0$. As the value of η_t is necessary for both steps of the two-step estimation routine, recovery of the structural parameters could be done using an iterative procedure between the first-stage conditional choice probability estimation and the second-stage utility parameter estimation. For any given value of η_t and s_t, one could recover a first-stage estimate of $FV'(s_t, x_t)$ and, conditional upon $\widehat{FV'}(s_t, x_t)$, the likelihood associated with the second-stage regression. The residual from the second-stage regression would itself be an estimate of η_t, which could be used to re-estimate $\widehat{FV'}(s_t, x_t)$ in the first stage and re-recover an updated $\widehat{\eta}$ in the second stage. This iterative procedure would continue until the estimates of η_t converge and would be required at each iteration of the likelihood.

3.4 Monte Carlo Evidence

In the appendix, we provide a simple, parameterized version of the model and illustrate its small-sample properties using Monte Carlo experiments. We compare estimates of the model using both the static and the dynamic estimators and the Monte Carlo results show that the static-model estimators return biased estimates of the marginal willingness to pay and that dynamic-model estimators show very little evidence of finite-sample bias.

4 Data

For this analysis, we employ a two-sided panel dataset describing both repeat sales of houses and repeat purchase decisions of buyers. The data cover five counties in the Bay Area of California (Alameda, Contra Costa, Marin, San Mateo, and Santa Clara) over the period 1990 to 2008.

The real estate transactions data were purchased from Dataquick and include dates, prices, loan amounts, and buyers', sellers', and lenders' names for all transactions. In addition, the data for the final observed transaction for each house include characteristics such as exact street address, square footage, year built, lot size, number of rooms, number of bathrooms, and number of units in the building.

The process of cleaning the data involves a number of cuts. Many of these are made in order to deal with the fact that we only see housing characteristics at the time of the last sale, but we need to use housing characteristics from all sales as controls in our hedonic price regressions. We therefore seek to eliminate any observations that reflect major housing improvement or degradation. First, to control for land sales or re-builds, we drop all transactions where "year built" is missing or with a transaction date that is prior to "year built". Second, in order to control for property improvements (*e.g.*, an updated kitchen) or degradations (e.g., water damage) that do not present as re-builds, we drop any house that ever appreciates or depreciates in excess of 50 percentage points of the county-year mean price change. We also drop any house that moves more than 40 percentile points between consecutive sales in the county-year distribution. Additionally, we drop transactions where the price is missing, negative, or zero. After using the consumer price index to convert all transaction prices into 2000 dollars, we drop one percent of observations from each tail to minimize the effect of outliers. Finally, as we merge-in data describing air quality and local crime rates using each property's geographic coordinates, we drop properties where latitude and longitude are missing. The publicly-available air quality data come from the California Air Resources Board.[17] The crime data are discussed below.

Summary statistics for the housing panel are given in Table 1. The sample is comprised of 541,415 transactions comprised of 359,151 distinct properties.

Table 1: Property Transactions Summary Statistics

N = 541,415					
Variable	Mean	Median	Std. Dev.	Min.	Max.
Year	1998.98	1999	5.12	1990	2008
House Age	32.03	31	20.52	0	147
House Square Footage	1,687.00	1545	662.12	160	9,130
Lot Square Footage	7,175.10	6,000	8,033.60	0	130,680
Number of Rooms	6.68	7	2.33	0	15
Number of Bedrooms	3.25	3	0.89	0	10
Number of Bathrooms	2.13	2	0.69	1	10
Ozone (one-hour peak annual concentration)	0.0638	0.0633	0.0087	0.0438	0.0860
PM10 (average annual concentration)	70.97	66.41	21.57	35.88	150.97
Property Crime Rate (per 100,000 residents)	1,608.10	1,392.80	783.28	245.70	11,522.00
Violent Crime Rate (per 100,000 residents)	379.63	322.65	262.88	12.82	3,834.10

[17]The air quality data are organized by monitor. There are 19 individual monitors in the six counties of interest. For each house, we generate a distance-weighted measure following the technique described for the crime data.

Crime statistics come from the RAND California data base, are organized by "city", and are measured as incidents per 100,000 residents. The data describe annual property and violent crime rates for each of the 80 cities within the San Francisco Metropolitan area. Figure 1 illustrates the location of these city centroids.

Figure 1: Cities within the San Francisco Metropolitan Area

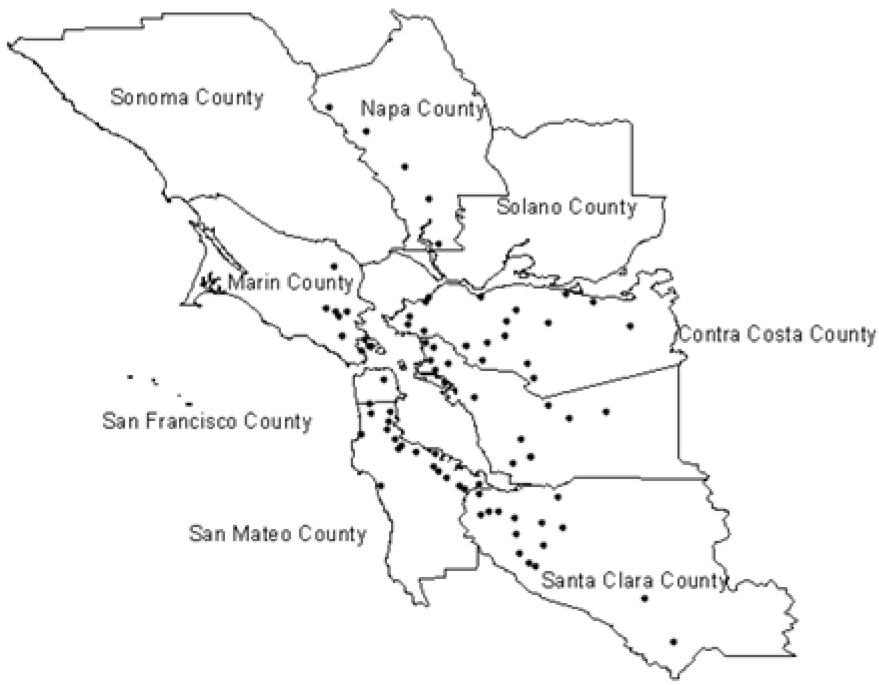

In this dataset, property crime is defined as "crimes against property, including burglary and motor vehicle theft", while violent crime is defined as "crimes against people, including homicide, forcible rape, robbery, and aggravated assault." Crime rates are imputed for each house in our data set using an inverse-distance weighted average of the crime rate in each city. Specifically, we weight the contribution of each city by the inverse of distance-squared, computing distance using the "great circle" calculation. We include the property crime rate as a control in our hedonic estimation and focus attention on violent crimes in our valuation exercise, as these crimes are less likely to be subject to systematic under-reporting. [Gibbons

(2004)]

Figures 2 and 3 illustrate the cross-sectional distribution of violent crime rates and the time-series of violent crime rates, respectively. There is a noticeable downward trend in violent crime, consistent with the decrease experienced by most of the US over this period.

Figure 2: Pooled Cross-Sectional Variation in Violent Crimes per 100,000 Residents

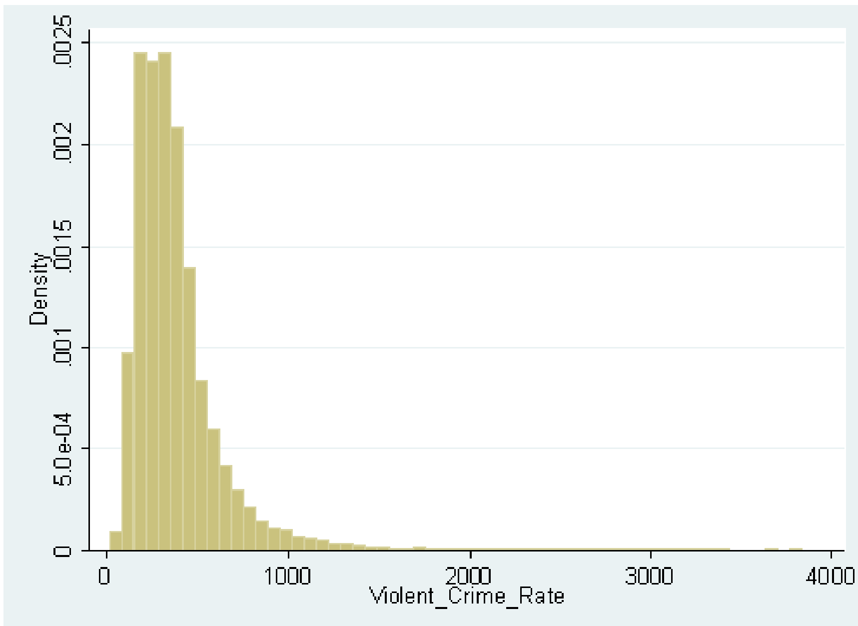

The air pollution measures come from the California Air Resources Board and are reported at the level of the air quality monitor, of which there are 21 located within the San Francisco Metropolitan area. The measure of air pollution that we use is the annual average concentration of PM10, or the concentration of particles with a diameter of 10 micrometers or less. This concentration is reported in micrograms per cubic meter of air. A measure of air pollution is generated for each house following the imputation described for crime.

In the implementation of our model, we also employ data following individual buyers through time. To create this panel, we first identify and match (through time) the "buyers" listed in the transaction records data described above. To accomplish this, we use the al-

Figure 3: Pooled Time-Series Variation in Violent Crimes per 100,000 Residents

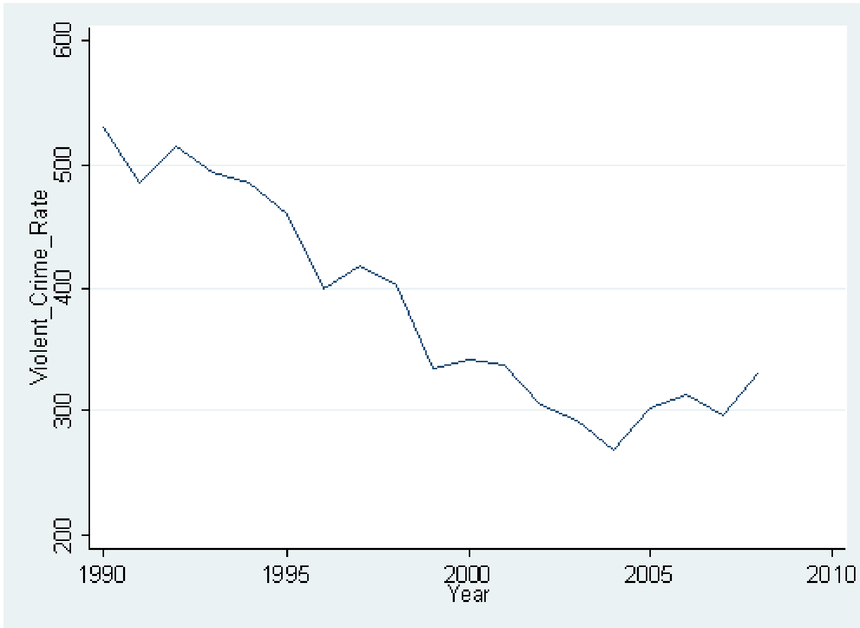

gorithm developed in Bayer, McMillan, Murphy, and Timmins (2016). Secondly, using the common variables of date, Census tract,[18] loan value, and lender, we merge-in data describing household race and income from the Home Mortgage Disclosure Act dataset (available for all households taking out a mortgage).

We successfully match approximately 75% of individuals in the transactions sample to the HMDA sample. Based on the algorithm for tracking households through time, we keep only those households observed to purchase three or fewer times during the sample period. Finally, we drop households where race or income are missing and households with income less than $25,000 or more than $500,000 income (in 2000 dollars). Note that this accounts for less than two percent of the remaining sample.

This yields a final sample of 372,334 observations comprised of 232,710 households.

[18]Census tracts are small, relatively homogenous geographic units defined by the Census Bureau. They contain between 2,500 and 8,000 individuals on average and vary in geographic size according to population density. [U.S. Census Bureau]

Summary statistics for the household panel are given in Table 2.

Table 2: Household Summary Statistics

N = 372,334					
Variable	Mean	Median	Std. Dev.	Min.	Max.
Income (year 2000 dollars)	118,824.80	102,000	68,110.10	25,000	500,000
White	0.58	1	0.49	0	1
Black	0.03	0	0.07	0	1
Asian	0.26	0	0.44	0	1
Hispanic	0.12	0	0.33	0	1

5 Results

5.1 Empirical Specification

We treat each county in the Bay Area as a separate market, denoted k, and allow the parameters of the rent function to vary by market.[19] Our primary amenities of interest, which we denote by x, are violent crime and air pollution (as measured by annual average concentration of PM10). In addition to our amenities of interest, we include as controls the additional amenities of property crime, property age, square footage, lot size, number of rooms, and full sets of Census-tract and year fixed effects and denote these controls h. The rental price function is given by:

$$log(r_{i,t}) = \gamma^r_{0,k} + x'_t\gamma^r_{1,k} + x'_t\gamma^r_{2,k}x_t + h'_t\gamma^r_3 + e_t$$

$r_{i,t}$ is the household's observed housing rent in the data. In practice, we follow the literature and define the yearly rental equivalent for owners as 0.075∗(property value). The deterministic

[19]We use data from the following five counties: Alameda, Contra Costa, Marin, San Mateo, and Santa Clara.

component of movings costs is given by $MC = z'_{i,t}\gamma^M$.

The direct component of utility is given by

$$u(x_t, \eta_t; \alpha) = \alpha_0 + x'_t(\alpha_{1,i} + \eta_t)$$

where $\eta \sim N(0, \Sigma_\eta)$ and $\alpha_{1,i} = z'_{i,t}\alpha_1$. The vector of household attributes, $z_{i,t}$, includes a constant, race, income, and year.

The transition probabilities of the j^{th} element of the amenity vector x are given by:

$$x^e_{j,t+1} = \rho_{0,j,k} + \rho_{1,j,k}x_t + \rho_{2,j,k}t + \nu_{j,t+1}$$

Letting $r'(x_{i,t})$ and $FV'(s_{i,t}, x_{i,t})$ equal the vector of derivatives the rental function and the future value with respect to each element of x, the first-order condition for the household's optimal choice of x is given by:[20]

$$\widehat{r}'(x_{i,t}) = z'_i\alpha_1 - \beta\sigma_\epsilon\widehat{FV}'(s_{i,t}, x_{i,t}) + \eta_{i,t} \tag{13}$$

where "hats" denote variables estimated in the first stages of estimation.[21] Using (13), to solve for $\eta_{i,t}$, we can form the likelihood of observing the continuous decision variable, x. The likelihood contribution of observation i is given by:[22]

$$L^x(\alpha, \Sigma) = (2\pi)^{-J/2}|\Sigma|^{-1/2}exp\{-\frac{\eta'_{i,t}\Sigma^{-1}\eta_{i,t}}{2}\}|\frac{\partial\eta_{i,t}}{\partial x_{i,t}}| \tag{14}$$

[20]For example the j^{th} element of $r'(x)$ is given by $\frac{\partial r(x)}{\partial x_j}$.

[21]Note that if we were to set $\beta = 0$, we would be faced with the standard (static) Rosen second-stage regression equation.

[22]This is the likelihood contribution of a household assuming that they moved and made a continuous choice. If they did not move, the endowment must have been consumed and the likelihood contribution is effectively one.

The marginal likelihood contribution of the discrete choice to form the likelihood is based on the probability of an agent choosing to move in any period. This probability is determined by the utility difference:

$$v_1(s_t) - v_0(s_t) = \big(x_t^*(s_t) - x_t^e\big)'\alpha_{1,i} - \big(\hat{r}(x_t^*(s_t)) - \hat{r}(x_t^e)\big) - z_{i,t}'\gamma^M +$$
$$\beta\sigma_\epsilon\big(E[-log(P_1(s_{t+1}))|s_t, d_t^m = 1, d_t^x = x_t^*(s_t)] - E[-log(P_1(s_{t+1}))|s_t, d_t^m = 0, d_t^x = x_t^e]\big) \quad (15)$$

The probability of moving is given by:

$$P_1(s_t) = \frac{e^{(v_1(s_t)-v_0(s_t))/\sigma_\epsilon}}{1 + e^{(v_1(s_t)-v_0(s_t))/\sigma_\epsilon}} \quad (16)$$

and the likelihood contribution of observation i is given by:[23]

$$L^m(\alpha, \Sigma, \gamma^M) = P_1(s_t)^{1(d_t^m=1)}(1 - P_1(s_t))^{1(d_t^m=0)} \quad (17)$$

Finally, to solve for both the discrete- and continuous-choice likelihoods, we must also estimate the reduced-form policy function $\hat{P}_1(s_t)$. We accomplish this by estimating the probability of reoptimizing with a flexible Logit separately for each market, k.

5.2 Results

For these results, we make the assumption that $\eta_{i,t}$ is observed after the household decides whether to re-optimize, but before the household decides how much crime to consume. We also only use the likelihood contribution associated with the continuous choice. We derive households' willingness-to-pay to avoid both violent crime and air pollution in the Bay Area

[23] An advantage of using information from the discrete choice is that it facilitates estimating the moving cost parameter, MC. Using only the continuous-choice, hedonic approach *controls* for moving costs, however, it does not recover an estimate of these costs.

of California. The estimation results from both the dynamic model and the myopic model are presented in Table 3.

Table 3: Estimation Results

	Dynamic	Static
VC - mean($z'_{i,t}\alpha_1$)	-10.07	-9.24
VC - α_{Income}	-0.03	-0.04
VC - σ_η	6.74	5.61
PM10 - mean($z'_{i,t}\alpha_1$)	-53.38	-37.29
PM10 - α_{Income}	-0.20	-0.21
PM10 - σ_η	37.57	35.50
σ_ϵ	16,925	–
ρ	0.417	-0.13
N	357,382	357,382

We find that the average household dislikes violent crime (mean($z'_{i,t}\alpha_1$) < 0)) and is willing to pay \$10.07 per year to avoid one additional crime per 100,000 residents. This translates to a willingness-to-pay of \$382.29 per year to reduce total violent crime by ten percent at the average level of violent crime (379.63 per 100,000 residents). We find that an additional \$1,000 in income increases willingness-to-pay by \$0.03, all else equal. This translates to an income elasticity of 0.36 calculated at the mean income of $\$118,824$ (in 2000 dollars).

The average household is willing to pay \$53.38 per year to avoid a one μ/m^3 increase in the local concentration of PM10. This translates to a willingness-to-pay of \$378.84 per year to reduce total PM10 by ten percent at the average level of PM10 (70.97 μ/m^3). We find that an additional \$1,000 in income increases willingness-to-pay by \$0.20, all else equal. This translates to an income elasticity of 0.46.

The standard deviation of the idiosyncratic taste parameters, the two σ_ηs, implies that tastes for crime and pollution are quite variable. The correlation coefficient, ρ, implies that these tastes are quite correlated, however.

The coefficient on the future value term, σ_ϵ, implies a reasonable standard deviation of moving costs of $\$30,699$.

In the myopic model, we find the average household is only willing to pay \$9.24 per year to avoid an additional violent crime per 100,000 residents and \$37.29 per year to avoid a one μ/m^3 increase in the local concentration of PM10. These figures represent an 8.24 and a 30.14 percent downward bias, respectively, when compared with the estimates from the dynamic model.[24]

6 Conclusion

The property value hedonic model, based on Rosen's 1974 seminal paper, has long been considered the "workhorse" model of amenity valuation. Derived using the household's first-order conditions associated with choosing where to live, the model is both intuitive and straightforward to estimate. However, changing residence is clearly a costly undertaking; real estate fees alone usually stand at six percent of a property's sales price and households typically have strong emotional ties to both their home and neighborhood. In the face of such costs, it becomes obvious that households behave dynamically when choosing where to live (and how much amenities to consume), yet the standard economic framework for modeling these decisions assumes that households are completely myopic. In this paper, we develop a model that allows for the fact that a change in the consumption of an amenity today affects not only current utility (through both increased consumption and increased price), but also the stream of future utility (through both future consumption levels and

[24]The test for whether the myopic willingness-to-pay is statistically different from the dynamic willingness-to-pay is simply the test for whether the coefficient on the future value term (σ_ϵ) is significantly different from zero.

future prices). We accomplish this while still retaining all of the intuition associated with the classic Rosen framework. Here, we simply redefine the familiar optimality condition to require that the marginal increase in current price is equal to the marginal benefit of the increase in amenity consumption this period *plus* the associated change in future utility flow. Finally, by presenting the household's problem as a two-part, discrete-continuous decision, we are able to take advantage of recent advances in the estimation of this class of model. We show how the added computational burden (compared with the static model) is reduced to a simple, first-stage logit probability estimation.

References

Albouy, D., W. Graf, R. Kellogg, and H. Wolff (2016): "Climate Amenities, Climate Change, and American Quality of Life," *Journal of the Association of Environmental and Resource Economists*, 3(1), 205–246.

Arcidiacono, P., and R. A. Miller (2011): "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity," *Econometrica*, 79(6), 1823–1867.

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

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

Bayer, P., R. McMillan, A. Murphy, and C. Timmins (2016): "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, 84(3), 893–942.

Bento, A., M. Freedman, and C. Lang (2015): "Who Benefits from Environmental

Regulation? Evidence from the Clean Air Act Amendments," *Review of Economics and Statistics*, 97(3), 610–622.

Bishop, K. C. (2012): "A Dynamic Model of Location Choice and Hedonic Valuation," mimeo, Arizona State University.

Bishop, K. C., and A. Murphy (2011): "Estimating the Willingness-to-Pay to Avoid Violent Crime: A Dynamic Approach," *American Economic Review, Papers and Proceedings*, 101(3), 625–629.

Bishop, K. C., and A. D. Murphy (2019): "Valuing Time-Varying Attributes using the Hedonic Model: When is a Dynamic Approach Necessary?," *Review of Economics and Statistics*, 101(1), 134–145.

Bishop, K. C., and C. Timmins (2017): "Estimating the Marginal Willingness to Pay Function Without Instrumental Variables," mimeo, Arizona State University.

Black, S. E. (1999): "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics*, pp. 577–599.

Brown, J. N., and H. S. Rosen (1982): "On the Estimation of Structural Hedonic Price Models," *Econometrica*, pp. 765–768.

Caetano, G. (2016): "Neighborhood Sorting and the Value of Public School Quality," mimeo, University of Rochester.

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

Chernozhukov, V., A. Galichon, M. Henry, and B. Pass (2017): "Single market nonparametric identification of multi-attribute hedonic equilibrium models," *mimeo.*

Davis, L. W. (2004): "The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster," *American Economic Review*, pp. 1693–1704.

Davis, M., J. Gregeory, D. Hartley, and K. Tan (2017): "Neighborhood Choices, Neighborhood Effects, and Housing Vouchers," Working Paper.

Downes, T. A., and J. E. Zabel (2002): "The Impact of School Characteristics on House Prices: Chicago 1987–1991," *Journal of Urban Economics*, 52(1), 1–25.

Dubin, J. A., and D. L. McFadden (1984): "An econometric analysis of residential electric appliance holdings and consumption," *Econometrica*, 52(2), 345–362.

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

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

Gayer, T., J. T. Hamilton, and W. K. Viscusi (2000): "Private Values of Risk Trade-offs at Superfund Sites: Housing Market Evidence on Learning about Risk," *Review of Economics and Statistics*, 82(3), 439–451.

Gibbons, S., and S. Machin (2003): "Valuing English Primary Schools," *Journal of Urban Economics*, 53(2), 197–219.

Greenstone, M., and J. Gallagher (2008): "Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program," *The Quarterly Journal of Economics*, 123(3), 951–1003.

Hanemann, W. M. (1984): "Discrete/continuous models of consumer demand," *Econometrica*, 52(3), 541–561.

Heckman, J., R. Matzkin, and L. Nesheim (2005): "Estimation and Simulations of Hedonic Models," *Frontiers in Applied General Equilibrium.*

——— (2010): "Nonparametric identification and estimation of nonadditive hedonic models," *Econometrica*, 78(5), 1569–1591.

Hotz, V. J., and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497–529.

Kennan, J., and J. R. Walker (2011): "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 79(1), 211–251.

Kuminoff, N. V., and J. C. Pope (2014): "Do "capitalization effects" for public goods reveal the public's willingness to pay?," *International Economic Review*, 55(4), 1227–1250.

Lee, S., and J. Lin (2018): "Natural amenities, neighbourhood dynamics, and persistence in the spatial distribution of income," *The Review of Economic Studies*, 85(1), 663–694.

Mangum, K. (2017): "The Role of Housing in Carbon Emissions," *Andrew Young School of Policy Studies Research Paper Series*, (17-05).

Mastromonaco, R. (2015): "A Dynamic General Equilibrium Analysis of School Quality Improvements," mimeo, University of Oregon.

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

Murphy, A. (2018): "A Dynamic Model of Housing Supply," *American Economic Journal: Economic Policy*, 10(4), 243–67.

Nesheim, L. (2015): "Identification of multidimensional hedonic models," *mimeo.*

Paciorek, A. (2013): "Supply constraints and housing market dynamics," *Journal of Urban Economics*, 77, 11–26.

Palmquist, R. B. (1982): "Measuring Environmental Effects on Property Values without Hedonic Regressions," *Journal of Urban Economics*, 11(3), 333–347.

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

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.

Appendix

A.1 Monte Carlo Evidence

We only show Monte Carlo simulations for the most parsimonious specification of our model – one with a single amenity of interest and with no household demographic attributes. To simulate the data necessary for a Monte Carlo experiment of a richer model specification, we would need to solve the full model (just to create the dataset). Thus, we would face the (prohibitively high) computational burdens discussed in Section 3.3. It is important to note, however, that the *estimation* of a richer model would not face such computational constraints, given the estimation strategy presented in this paper.

For simplicity, let x be a scalar amenity (a good) and let the number of markets, J, be equal to two. The direct component of utility is given by:

$$u(x_t, \eta_t; \alpha) = \alpha_0 + (\alpha_1 + \eta_t)x_t$$

where $\eta \sim N(0, \sigma_\eta)$. We assume the model timing outlined in Section 3.3.1, with the shock to preferences, η_t, being observed after the decision to move has been made, but before the decision of how much x to consume has been made.

The rental price function, of which we allow the parameters to vary by market, is given by:

$$r(x_t; \gamma^r) = \gamma^r_{0,j} + \gamma^r_{1,j}x_t + \gamma^r_{2,j}x_t^2 + e_t$$

and the deterministic component of movings costs is given by the constant, MC. The exogenous transition probabilities of the amenity x are given by:

$$x^e_{t+1} = \rho_0 + \rho_1 x_t + \nu_{t+1}$$

where $\nu \sim N(0, \sigma_\nu)$.

Using this same basic structure, we recover the parameters using (i) the likelihood of observing the continuous choice, (ii) the likelihood of observing the discrete choice, and (iii) the joint likelihood of the two choices. We additionally present the results from two static approaches for comparison.

A.1.1 Monte Carlo Results

In Table A.1, we compare estimates of the model using both the static and the dynamic estimators. In Columns I and II, we present the results using the static hedonic approach and the static discrete-choice approach. The results of the dynamic hedonic estimator, the dynamic discrete-choice estimator, and joint-likelihood estimator are shown in Columns III, IV, and V, respectively. In each experiment, we set the number of draws to 500, the number of households to 1,000, and the number of time periods to 20. We fix and $\beta = 0.95$ and assume that it is know by the econometrician.

The Monte Carlo results show that the static models return biased estimates of the primary structural parameter of interest, α_1. The estimates of the parameters from all three dynamic estimators show very little evidence of finite-sample bias.[25] Among the dynamic estimators, the variance of the dynamic hedonic estimator shown in Column III is somewhat higher than that of the dynamic discrete-choice estimator shown in Column IV.[26] This variance differential is caused by two factors: first, the sample size is approximately four times larger for the dynamic discrete-choice estimator, as data describing the continuous choice of x_t is only available for households who, in fact, move (the unconditional probability of moving is 0.252), and, second, the estimation of the dynamic hedonic estimator involves estimating a non-parametric derivative which, while not affecting the consistency of the estimator, does add to the variance. That said, the simplicity and computational tractability of the dynamic hedonic estimator make it an attractive estimation approach.

Table A.1: Monte Carlo Results

		I	II	III	IV	V
	True Values	Static Hedonic	Static Discrete Choice	Dynamic Hedonic	Dynamic Discrete Choice	Dynamic Joint Likelihood
α_1	4	3.680	3.772	4.014	3.996	3.996
		(0.003)	(0.017)	(0.067)	(0.018)	(0.012)
σ_ϵ	1.2	–	–	1.278	1.196	1.206
		–	–	(0.271)	(0.034)	(0.051)
σ_η	0.25	0.230	–	0.252	0.335	0.250
		(0.002)	–	(0.012)	(0.374)	(0.007)
MC	3	–	2.466	–	3.177	3.012
		–	(0.041)	–	(0.376)	(0.088)

[25] σ_η is poorly identified in the dynamic discrete-choice estimator, as in this case it is only identified off functional form. Apart from, σ_η, the estimate with the largest finite-sample bias is σ_ϵ. The mean of σ_ϵ from the dynamic hedonic approach (Column III) is 1.278. The median estimate of 1.226 is much closer to the true parameter.

[26] The exception to this is the estimated variance of σ_η.