Valuing Time-Varying Attributes using the Hedonic Model: When is a Dynamic Approach Necessary?*

Kelly C. Bishop†
Arizona State University

Alvin D. Murphy‡
Arizona State University

February 2016

Abstract

We build on the intuitive (static) modeling framework of Rosen (1974) and specify a simple forward-looking model of location choice. We use this model, along with a series of insightful graphs, to describe the potential biases associated with the static approach and relate these biases to the time-series trend of the amenity of interest. We then derive an adjustment factor that allows the biased static estimates to be converted into estimates coming from a forward-looking model. Finally, we empirically motivate the use of this adjustment factor in an application of estimating the willingness-to-pay to avoid violent crime in California.

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

*We would like to thank Esteban Aucejo, Pat Bayer, Morris Davis, Dennis Epple, Jesse Gregory, Nick Kuminoff, Ralph Mastromonaco, Rob McMillan, Lucija Muehlenbachs, Steve Tadelis, as well as various conference and seminar participants for their helpful comments and suggestions.
†Kelly C. Bishop, Arizona State University, Box 879801, Tempe, AZ 85287. kelly.bishop@asu.edu
‡Alvin D. Murphy, Arizona State University, Box 879801, Tempe, AZ 85287. alvin.murphy@asu.edu

1 Introduction

The standard hedonic model, drawing on Rosen's classic 1974 paper, provides the workhorse empirical approach used to value local public and private goods. It is straightforward to estimate, usually involving a single least-squares regression of house prices on housing characteristics and neighborhood amenities; applying the model has become ever-more feasible with the increasing availability of detailed housing transactions data. Given its appeal, a myriad of hedonic valuation exercises have been featured in the literature, 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)), and many more.

An implicit assumption underlying the traditional model is that households are myopic – that is, they do not account for the fact that housing and neighborhood amenities are likely to be time-varying. In practice, though, given the significant costs associated with purchasing a house and with moving, it is unlikely that households would not consider future levels of local amenities when making their decisions. When households are forward-looking in this manner, the traditional model will yield biased estimates of willingness-to-pay in many cases, but not all, and the degree of bias will vary.

A recent literature has sought to quantify this bias empirically in specific applications, comparing results from the traditional model to those obtained using fully-dynamic models of location choice.[1] Yet the estimation of dynamic models comes with substantial computational costs, even when drawing on recent advances in the literature. Furthermore, fully dynamic models often require very rich data that may exceed the detail of existing data sets. Thus, it would be useful if applied researchers could determine in advance of a full-blown dynamic estimation whether the resulting benefits were likely to outweigh the significant computational and data costs involved.

In this paper, we provide a framework that allows such a pre-determination. Building on the static modeling framework of Rosen (1974), we specify a simple forward-looking model of location choice where households choose a residence based on the stream of associated utility flows for a fixed number of years. Using this framework, we characterize more fully the potential bias associated with the static approach and relate this bias to the time-series trend

[1] For recent papers that estimate dynamic models of location choice, see Kennan and Walker (2011), Bishop (2012), Bayer, McMillan, Murphy, and Timmins (2015), Bishop and Murphy (2015), Caetano (2015), Davis, Gregeory, Hartley, and Tan (2015), and Mastromonaco (2015).

of the amenity of interest. In addition, we illustrate an empirically-relevant example where the static model and the forward-looking model arrive at the same estimate of willingness to pay, despite the fact that the amenity in question is time-varying.

To understand the potential bias arising from a static model, it is worthwhile revisiting the intuitive identification strategy of the Rosen model which allows the researcher to recover estimates of marginal willingness to pay for an amenity from information about (i) the quantity of the amenity the household chooses to consume and (ii) the price schedule faced by the household. First, the parameters of the housing price function (and, at the same time, the parameters of the hedonic price gradient) are recovered through a regression of observed housing prices on amenities. Second, the econometrician is able to back out the implicit price of the amenities that each household actually paid using the family's observed consumption of the amenity and the parameters of the hedonic price gradient. The information provided by the first-order condition for utility maximization (i.e., that marginal cost will equal marginal benefit) allows the econometrician to arrive at the household's marginal willingness to pay for the amenity.

If, however, households are choosing where to live based on some average stream of future amenities and not solely based on a measure of current amenities, as is likely, then the traditional model will get both (i) and (ii) wrong, resulting in potentially biased estimates of willingness to pay. It is straightforward to see that by using the incorrect measure of quantity consumed, the static model will either under- or over-attribute the true quantity purchased. We refer to this as the *quantity effect*. Less obvious is what we refer to as the *price effect*: if the econometrician is recovering the implicit price of the amenity through the calculation of the price differential under quantity mismeasurement, the implicit price of the amenity will also be under- or over-stated. Using these quantity- and price-effect notions, we seek to describe more fully the potential bias associated with the static approach and relate this bias to the time-series properties of the amenity of interest.

More specifically, we show that in cases where the location-specific amenity levels are mean-reverting over time within the choice set, the traditional static model will typically understate willingness-to-pay (i.e., coefficients will be biased toward zero). For example, consider a household that purchases a house with a lower-than-mean amenity level in the current period. As a result of the mean-reversion over time, the average stream of the consumed amenity will be higher than the current level. In addition, the true implicit price of the amenity will be higher than that obtained using the static approach. In other words, the household is buying relatively more of the amenity at a relatively higher price and therefore the true willingness to pay is higher than is implied by the static model.[2] Analogously, the traditional model will typically overestimate willingness-to-pay (i.e., coefficients will be biased away from zero) in cases where the levels of the amenity are diverging across locations.[3] Intuitively, the size of the bias is determined by how quickly the amenity is diverging through time.

Following this quantity- and price-effect intuition, we present a result where the forward-looking model and the static model yield the same estimate of willingness to pay, even when the amenity is evolving rapidly through time. This arises when the amenity of interest is rising or falling without mean-reverting or mean-diverging, thereby causing the effect of the misspecified quantity to exactly offset the effect of the misspecified price in the static model. For example, if each house received a one-unit increase in the amenity, then amenity levels would be increasing over time and the static model would understate quantity and overstate price. We show that in the case of linear utility, these effects exactly offset one another and the static model yields unbiased estimates of willingness to pay.

We then show that an adjustment factor may be easily derived. This adjustment factor, which is based on the time-series properties of the amenity of interest, can be then be used to convert the biased static estimates into the estimates that one would have obtained using the forward-looking model. In the simplest case, this adjustment factor is a scalar which can be obtained from a simple ordinary least squares regression.

Finally, using a rich data set on housing transactions, we apply our adjustment factor to illustrate examples where the static model yields large biases and examples where the static model yields small biases. In particular, we use data from the San Francisco Bay Area to estimate the marginal willingness to pay to avoid crime. We calculate the bias separately by county and find that the static model produces a small bias in Alameda County, while producing large biases in both Marin and San Mateo counties. In fact, the static model understates willingness to pay to avoid crime by a factor of two in the latter cases. The heterogeneity across counties is driven by the fact that there is only a small amount of mean reversion in Alameda County, while crime mean-reverts quickly in Marin and San Mateo Counties. This geographic heterogeneity provides an empirical example which supports our assertions that, depending on the application, the bias generated by specifying a static model

[2] If the household started out consuming a high level of the amenity at present, the quantity and price effects would go in opposite directions. That is, relative to the static model, the household would be consuming less of the amenity but at a higher price. We will show that, in this case, the price effect dominates and the static model would underestimate willingness to pay.

[3] Consider again the household that purchases a house with a low level of an amenity at present. Now, both the consumption and the true implicit price of the amenity will be overstated by the static model. In other words, the household is buying relatively less of the amenity and paying a lower price and as such, the true willingness to pay is lower.

may be large or may be small and that it is straightforward to get a sense of this bias without estimating a fully-dynamic model.

The remainder of the paper is organized as follows: Section 2 describes the traditional static model as well as a simple forward-looking model of hedonic demand; Section 3 describes the bias induced by the misspecified model under various transitions of the amenity of interest and provides guidance when trying to answer the question, "when is the static model sufficient?"; Section 4 applies the framework to estimate the willingness to pay to avoid crime; and Section 5 concludes.

2 Model

In this section, we provide an overview of the traditional, Rosen-style static model, as well as a simple forward-looking model of willingness to pay.

2.1 The Traditional, Static Model of Willingness to Pay

We first consider the static model of willingness to pay for a house or neighborhood amenity. In this model, households maximize current utility with respect to their choice of amenity consumption.[4] We choose a simple specification of household utility where household i has an individual-specific preference parameter, α_i, describing their preference for consumption of the amenity of interest, x_i. The household also receives utility from the consumption of the numeraire good, C_i.

$$U(x_i) = \alpha_i x_i + C_i \tag{1}$$

For simplicity, we consider a model where utility is increasing at a linear rate in the amenity x.[5] Broadly speaking, the intuition developed here applies to non-linear specifications, which we present in the Appendix.

Households purchase x as part of the bundle of goods described by housing. Households must pay an annual user cost for housing, which we denote r_i. One could think of the annual

[4]Implicitly assumed in this modeling framework is the assumption of free mobility (or zero transaction costs). This means that households can freely reoptimize at the beginning of each period, so the problem of maximizing lifetime utility may be described as a series of independent, sequential decisions.

[5]While this is done to simplify the analysis, it also means that the identification issues discussed in Brown and Rosen (1982), Mendelsohn (1985), and Ekeland, Heckman, and Nesheim (2004), as well as the estimation issues discussed in Epple (1987) and Bartik (1987), do not apply here.

user cost of housing as capturing either a rent or mortgage payment.[6] The function that relates the level of amenity consumption, x, to the annual user cost of housing, r_i, is the housing price function, $r_i = r(x_i)$. Incorporating the household's budget constraint, i.e., that their numeraire consumption, C_i, is equal to income, I_i, minus annual user cost of housing, r_i, yields:

$$U(x_i) = \alpha_i x_i + I_i - r(x_i) \tag{2}$$

According to the assumptions underlying the static model, households only consider current levels of the amenity x and, therefore, maximize current utility with respect to their current choice of x.[7] Thus, the first-order condition for the optimal choice of x is given by:

$$U'(x_i) = \alpha_i - r'(x_i) = 0 \tag{3}$$

The first-order condition described by Equation 3 may then be used to solve for α_i, household i's marginal willingness to pay for amenity x.[8] In other words, at their chosen level of x consumption, household i's marginal utility of x will equal the implicit price of x, i.e., the marginal cost:

$$\alpha_i = r'(x)\Big|_{x=x_i^*} \tag{4}$$

which naturally suggests the (static) estimator,

$$\hat{\alpha}_i^s = \hat{r}'(x)\Big|_{x=x_i^*}, \tag{5}$$

for the per-annum willingness to pay for a one-unit increase in the amenity, x.

In addition to this analytical solution, the estimation framework may be described intuitively through a series of graphs. In a first stage, the parameters of the housing price function and, at the same time, the parameters of the implicit price function (hedonic price gradient), $r'(x)$, are recovered through a regression of annual user costs of housing on amenity levels. These relationships are shown in Figure 1.[9]

In a second stage, the econometrician is able to back-out the implicit price of x that

[6]The annual user cost of housing could also capture other costs of home ownership such as taxes, maintenance, and depreciation. See Poterba (1984) for a discussion of user cost.

[7]If x is time-varying, this assumption is analogous to an assumption regarding households' moving costs; in a world with zero moving costs, households may costlessly reoptimize in every period (so looking to the future yields no benefit). However, many papers, including Kennan and Walker (2011), find evidence of substantial moving costs.

[8]See Bajari and Benkard (2005) and Bajari and Kahn (2005) for an insightful discussions of the interpretation of α_i as a structural parameter of the utility function, rather than simply a local estimate of marginal

Figure 1: Graphical Representation of the First-Stage Estimation

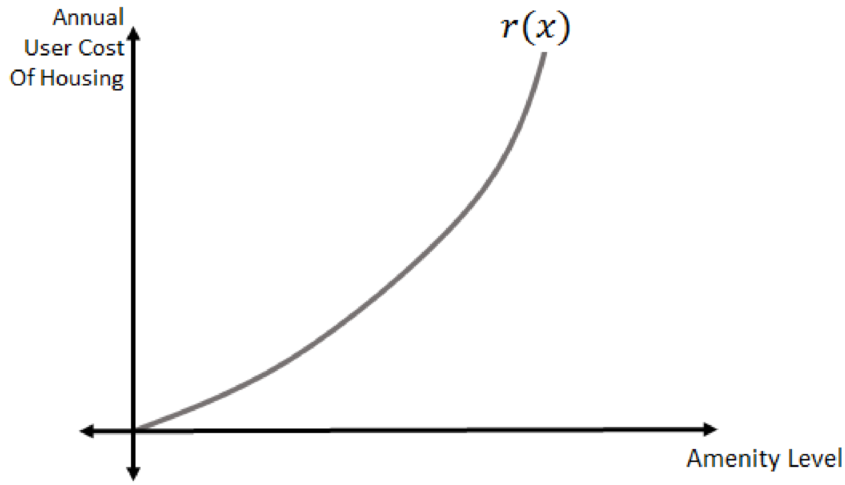

(a) The Housing Price Function, $r(x)$

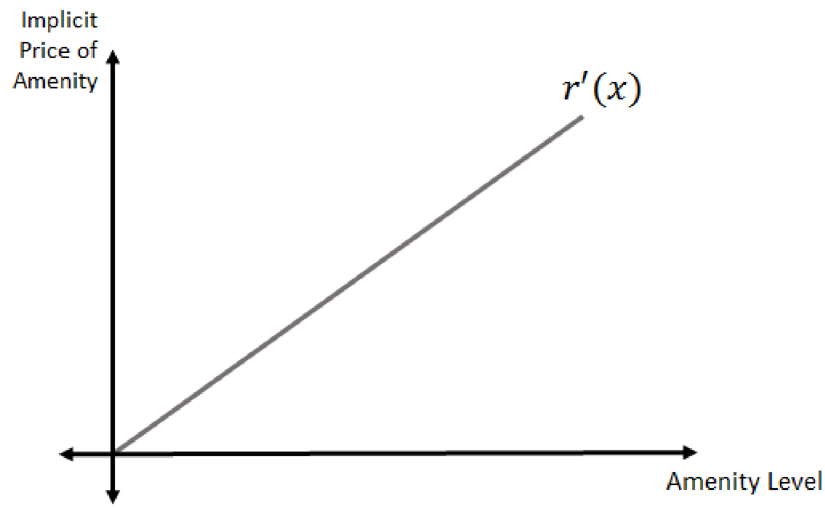

(b) The Implicit Price Function, $r'(x) = \frac{\partial r(x)}{\partial x}$

each family actually paid, using the family's observed consumption of x and the implicit price function. The information provided by the first-order condition for utility maximization allows the econometrician to equate this implicit price to arrive at the preference parameter,

willingness to pay at the point of consumption.

[9]In the figures, we illustrate a quadratic housing price function (linear implicit price function) and an amenity that may be considered a "good". However, the intuition and results hold for any form of the housing price function.

α_i. This second stage is depicted in Figure 2.

It is important to note that this static model will only return unbiased estimates of α_i when either moving is costless or when amenity levels are fixed through time. However, in any realistic application, a household would face positive moving costs and time-varying amenities. Thus, we describe a forward-looking model in the next section.

Figure 2: Graphical Representation of the Second-Stage Estimation

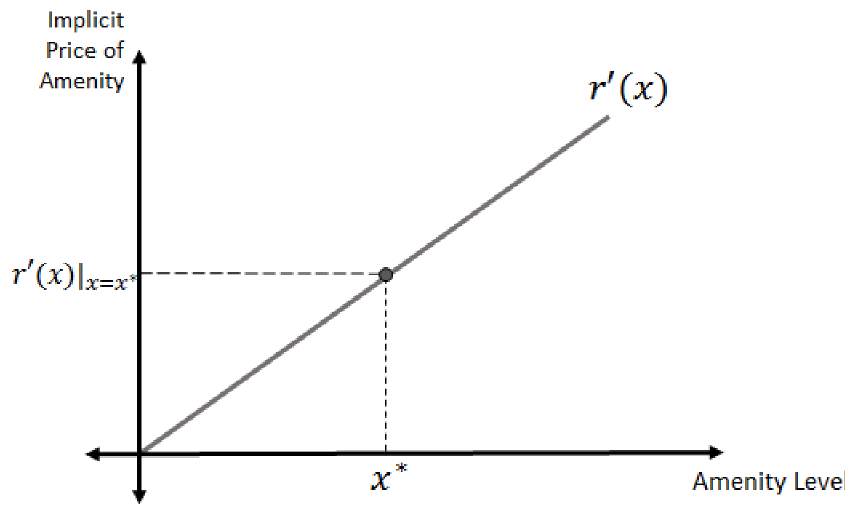

(a) Household i's Observed Amenity Consumption

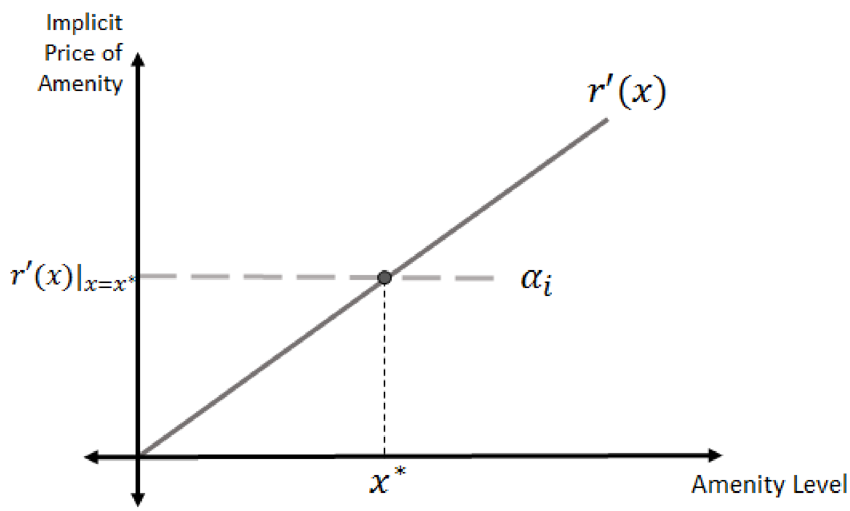

(b) Household i's Marginal Willingness to Pay

2.2 A Simple, Forward-Looking Model of Willingness to Pay

We now move to a forward-looking framework where households maximize the discounted sum of annual utility flows with respect to their current choice of x. Our goal is to specify a model that captures the key determinants of forward-looking behavior but is still simple enough to retain analytical tractability. To do this, we abstract away from some of the finer details of dynamic behavior that would substantially complicate the analysis and preclude an analytical decomposition of the bias. In the empirical application of Section 4, we compare results from the simple forward-looking model presented here with those found in the fully-dynamic model of Bishop and Murphy (2011).

We assume that households choose a residence based on the stream of associated utility flows for the next T years.[10] This is akin to assuming prohibitively high moving costs for the next T years. For simplicity, we assume households cannot reoptimize within the period of T years and we abstract away from any considerations about the post-T utility. In the specification laid out here, we do not consider future reoptimization in order to simplify the problem, yet we retain the primary insights and intuition of a fully-dynamic model – that is, households know that amenities are time varying and that their choice of amenity today will influence the amount of the amenity they consume in subsequent periods.[11]

The housing price function still maps the consumption of the amenity into the annual user cost of housing. As this price is determined in the current period, t, it is a function of current amenity levels and denoted $r(x_{i,t})$. For homeowners, who are our group of interest, it is natural to think of this annualized user cost of housing as a mortgage payment: determined at the time of sale, it is a function of amenity levels in the period in which a household buys.[12]

The amenity of interest, x, is evolving through time and households form expectations about future levels of x. Denoting the current period as t (and current choice of amenity level as $x_{i,t}$) we can write the discounted sum of annual utility flows over the next T years (i.e.,

[10] In the application, we set T to seven years, which is approximately the median household tenure in the United States over this period. See Section 4 for a discussion.

[11] In a fully-dynamic model, households would also maximize the discounted sum of annual utility flows (*i.e.*, lifetime utility), but would face positive, yet feasible, moving costs in each period. Households would then account not only for future utility flows, but for possible future endogenous reoptimization from their current choice.

[12] A more general model would allow the user cost to vary over time, which would complicate the model and analysis. As discussed above, the goal of the paper is to derive simple analytical results for the simplest model that still captures the key component of dynamic behavior. In that spirit, it is natural to restrict the user cost to be time invariant, but to allow utility from the amenity to vary over time.

the value function) as:

$$v(x_{i,t}) = E[\sum_{s=1}^{T} \beta^{s-1}(\alpha_i x_{i,t+s-1} + I_i - r(x_{i,t}))|x_{i,t}] \tag{6}$$

which can be rewritten as:

$$v(x_{i,t}) = \alpha_i \sum_{s=1}^{T} \beta^{s-1} E[x_{i,t+s-1}|x_{i,t}] + \sum_{s=1}^{T} \beta^{s-1} I_i - \sum_{s=1}^{T} \beta^{s-1} r(x_{i,t}) \tag{7}$$

For exposition purposes, we define a measure of expected average x consumption over the horizon T with the following weighted average:

$$\bar{x}_{i,t} = \bar{x}(x_{i,t}) = \frac{\sum_{s=1}^{T} \beta^{s-1} E[x_{i,t+s-1}|x_{i,t}]}{\sum_{s=1}^{T} \beta^{s-1}}$$

We also define the function $\tilde{r}(\bar{x})$ which maps the expected average stream of amenity flows into the annual user cost of housing, $r_{i,t}$:

$$r_{i,t} = \tilde{r}(\bar{x}(x_{i,t})) = r(x_{i,t}) \tag{8}$$

To make this concrete, consider a given house has an annual user cost of \$10,000, then $r_{i,t} = 10,000$. If this house has a current level of amenities equal to 90 and an expected average level of amenities of 115, then $\tilde{r}(115) = r(90) = 10,000$.

Defining $\tilde{v}(\bar{x})$ analogously (i.e., that $\tilde{v}(\bar{x}(x_{i,t})) = v(x_{i,t})$), allows us to rewrite Equation (7) in terms of this average x:

$$\tilde{v}(\bar{x}_{i,t}) = \alpha_i \sum_{s=1}^{T} \beta^{s-1} \bar{x}_{i,t} + \sum_{s=1}^{T} \beta^{s-1} I_i - \sum_{s=1}^{T} \beta^{s-1} \tilde{r}(\bar{x}_{i,t}) \tag{9}$$

The household's problem is then equivalent to choosing $\bar{x}_{i,t}$ to maximize $\tilde{v}(\bar{x}_{i,t})$, yielding the first-order condition:[13]

$$\tilde{v}'(\bar{x}_{i,t}) = \alpha_i \sum_{s=1}^{T} \beta^{s-1} - \sum_{s=1}^{T} \beta^{s-1} \tilde{r}'(\bar{x}_{i,t}) = 0 \tag{10}$$

[13] If one were to work with $x_{i,t}$ instead of $\bar{x}_{i,t}$, the first-order condition would be given by: $\partial \tilde{v}(\bar{x}_{i,t})/\partial x_{i,t} = \alpha_i \partial \bar{x}_{i,t}/\partial x_{i,t} - \tilde{r}'(\bar{x}_{i,t}) \partial \bar{x}_{i,t}/\partial x_{i,t} = 0$ which is equivalent to Equation 10.

The first-order condition described by Equation 10 may then be used to solve for household i's marginal willingness to pay for amenity x. In other words, at their chosen level of $\bar{x}$ consumption, household i's individual-specific preference parameter, α_i, can be recovered as:

$$\alpha_i = \tilde{r}'(\bar{x})\Big|_{\bar{x}=\bar{x}^*_{i,t}} \tag{11}$$

which naturally suggests the (forward-looking) estimator,

$$\hat{\alpha}_i^f = \hat{\tilde{r}}'(\bar{x})\Big|_{\bar{x}=\bar{x}^*_{i,t}}, \tag{12}$$

for the per-annum willingness to pay for a one-unit increase in the amenity, x.

When compared with the analogous solution from the static model (described by Equation (5)), Equation (12) highlights the two effects that we previously referred to as the price effect and the quantity effect. The price effect is captured by the use of $\tilde{r}(\cdot)$ rather than $r(\cdot)$. The quantity effect is captured by the fact that we evaluate the function at $\bar{x}_{i,t}$ rather than $x_{i,t}$. In the following section, we discuss the bias induced by each of these effects and show an interesting result where these two effects cancel one another out.

Note that graphically the recovery of α_i for the forward-looking model appears similar to that of the static model depicted in Figures 1 and 2, but defined in $(\tilde{r}'(\bar{x}), \bar{x})$ space.

3 Understanding and Predicting the Bias

When the amenity of interest is time-varying and reopimization is not without cost, estimates of willingness to pay recovered using the static model may be biased. In this section, we provide a detailed decomposition of the bias by relating it to the time-series properties of the amenity of interest, x. We do this using a series of intuitive graphs and by discussing the mathematical difference between the estimate of marginal willingness to pay recovered from the static model (Equation 5) and the estimate of marginal willingness to pay from the forward-looking model (Equation 12).

The transition properties of the amenity of interest will determine the sign and size of the bias and will, therefore, determine when the estimation of the dynamic model is most warranted. When considering the time trend of the amenity, it is sufficient to describe two key features: (i) what is the overall trend over the next T years, i.e, is the expected average amenity level higher or lower than current amenity level? and (ii) is the amenity level mean reverting or mean diverging? In the remainder of this section, we walk through the various potential paths of housing amenities and discuss their impacts on willingness-to-pay estimates.

3.1 The Amenity is Rising or Falling Through Time

We first consider the case where the amenity of interest, x, is either simply rising or falling through time in a manner that preserves the variance in x across locations. In other words, the relationship between x_{it} and $\bar{x}_{it}$ can be expressed as:

$$\bar{x}(x_{i,t}) = \phi + x_{i,t} \tag{13}$$

In this example, ϕ can be either positive or negative and we put no restrictions on its magnitude. When $\phi > 0$, the amenity is rising over the T-year horizon and the average future amenity level, $\bar{x}_{i,t}$ will be higher than the current amenity level, $x_{i,t}$. This would be the case if the amenity were local expenditure on public schools and all schools received the same dollar increase in budget. Alternatively, when $\phi < 0$, the amenity is falling through time. In either case, this ensures that a change in current $x_{i,t}$ produces a one-for-one change in average future amenity consumption, $\bar{x}$. An alternative way of stating this is that for any two choices of current x, denoted x_a and x_b: $|\bar{x}(x_a) - \bar{x}(x_b)| = |x_a - x_b|$.

Graphically, a uniform increase in x ($\phi > 0$) is represented in Figure 3. In Figure 3a, one can see that the increase in x results in a forward-looking price function that lies (in a parallel manner) to the right of the static price function. In other words, for any given level of housing expenditure, the associated average amenity level, $\bar{x}_{i,t}$, is higher than the current amenity level, $x_{i,t}$. Correspondingly, the forward-looking implicit price function, which is depicted in Figure 3b, lies (in a parallel manner) to the right of the static implicit function. In other words, the implicit price of the amenity is lower than the static model would imply.

As can be noted in Figure 3b, the quantity effect and the price effect work in opposite directions and *exactly* offset one another; there is no bias associated with the static modeling framework, even when facing an amenity with potentially strong time-trends (as long as the time-trend affects all locations of the choice set uniformly). Even with the static modeling framework, the econometrician will recover the true estimate of willingness to pay.

Figure 3: The Amenity is Rising Through Time

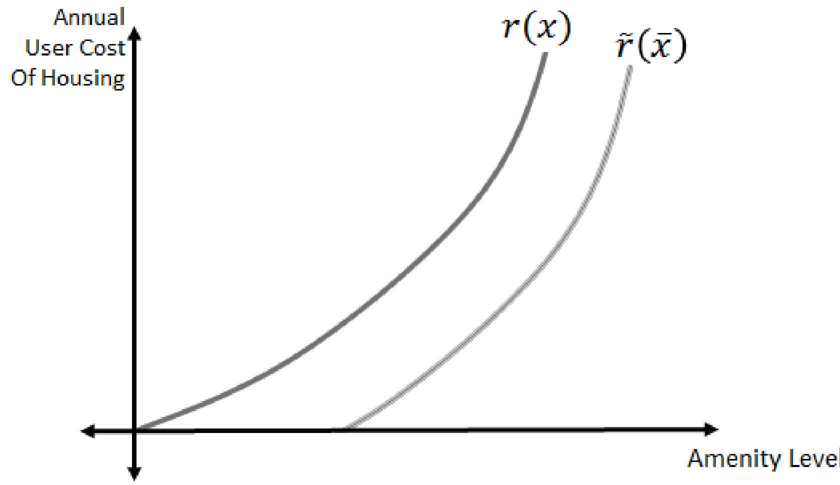

(a) The Housing Price Functions

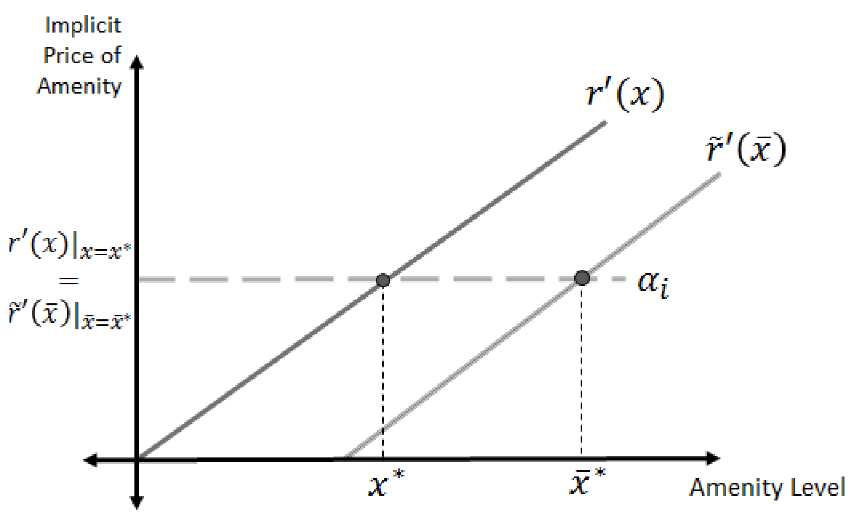

(b) The Implicit Price Functions and Household i's MWTP

Graphically, the quantity and price effects are labeled in Figure 4.[14] As Figure 4 shows, the household's average consumption of x is higher than than the static model would imply, leading to a positive quantity effect. The implicit price of the amenity is lower than the static model would imply, leading to a negative price effect.

[14]The order is which we illustrate the quantity effect and the price effect is arbitrary. That is, we could have alternatively defined the quantity effect as moving from $\tilde{r}'(x^*)$ to $\tilde{r}'(\bar{x}^*)$ and price effect as moving from $r'(x^*)$ to $\tilde{r}'(x^*)$. In either case, the net result is the same.

Figure 4: The Competing Effects When x is Increasing

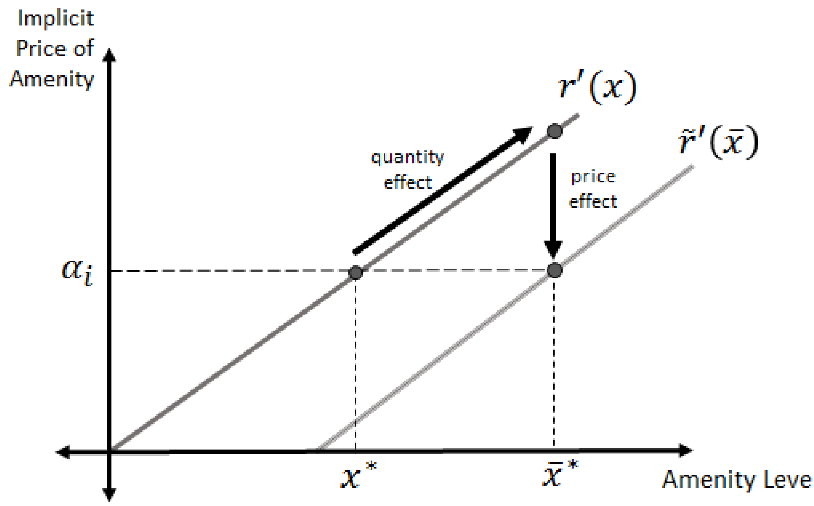

We additionally illustrate a simple decrease in the amenity in Figure 5.

Figure 5: The Competing Effects When x is Decreasing

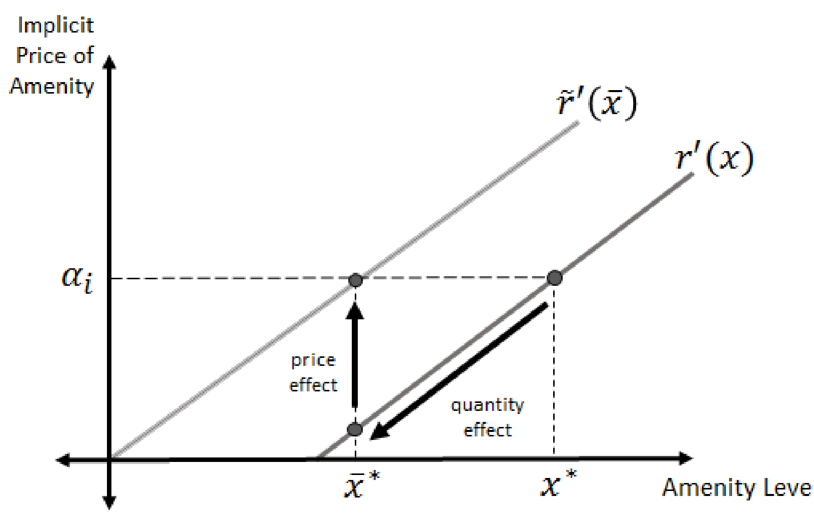

This figure again shows the quantity effect and the price effect working in opposite directions and exactly offsetting one another. In this case, the household's average consumption of the amenity is lower than the static model would predict, leading to a negative quantity effect. However, the implicit price of the amenity is higher than the static model would predict, leading to a positive price effect.

Analytically, we can derive an expression for the difference between that static and forward-looking estimates. Using the definition that appears in Equation 8, that $\tilde{r}(\bar{x}(x_{i,t})) = r(x_{i,t})$, and differentiating with respect to $x_{i,t}$ yields:

$$\tilde{r}'(\bar{x}_{i,t})\frac{\partial \bar{x}_{i,t}}{\partial x_{i,t}} = r'(x_{i,t}) \tag{14}$$

It follows that:

$$\underbrace{\tilde{r}'(\bar{x})\Big|_{\bar{x}=\bar{x}^*_{i,t}}}_{\text{MWTP forward-looking model}} = \frac{1}{\frac{\partial \bar{x}_{i,t}}{\partial x_{i,t}}} \underbrace{r'(x)\Big|_{x=x^*_{i,t}}}_{\text{MWTP static model}} \tag{15}$$

It can be easily seen in Equation 15 that a single, scaling term captures the difference between the marginal willingness-to-pay estimate from the static model and that from the forward-looking model; the quantity and price effects are both captured by a single term which fully describes the bias. This term, $\partial \bar{x}_{i,t}/\partial x_{i,t}$, explicitly shows that both the sign and size of the bias are driven by how a household's current choice of amenity level affects the average future stream of amenity levels. If one were to recover this term in a separate first stage, it presents a simple and practical adjustment to estimates from the static model: divide through by $\partial \bar{x}_{i,t}/\partial x_{i,t}$. This expression will be used in our empirical application to convert static estimates into the estimates one would have obtained using the forward-looking model.

Returning to the simply rising- (or falling-) amenity case where $\bar{x}(x_{i,t}) = \phi + x_{i,t}$, it is clear that:

$$\partial \bar{x}_{i,t}/\partial x_{i,t} = 1 \tag{16}$$

Based on Equation 15, this implies that the bias term, $\partial \bar{x}_{i,t}/\partial x_{i,t}$, drops out and the willingness to pay derived using the static model is identical to that derived using the dynamic one:

$$\tilde{r}'(\bar{x})\Big|_{\bar{x}=\bar{x}^*_{i,t}} = r'(x)\Big|_{x=x^*_{i,t}} \tag{17}$$

3.2 The Amenity is Mean Reverting or Mean Diverging

We now consider changes in the amenity of interest that are not uniform across the locations of the choice set. In other words, we consider cases where amenity levels are either mean reverting or mean diverging over the T-year horizon. The simplest case of mean reversion would arise if shocks to amenity levels arrive through time and these shocks decay. More complicated cases of mean reversion could be the result of targeted policy; for example, in the case of school quality, resources may be diverted to districts with the lowest performance in the prior period. Another example would be tipping points. Tipping points in neighborhood racial composition would be an example of an amenity displaying mean divergence; neighborhoods with minority levels above some tipping point will experience more in-migration of minorities, while neighborhoods with minority levels below some tipping point will experience more out-migration of minorities. For our purposes, however, we do not distinguish between the underlying causes of mean reversion or mean divergence.

We now consider a relationship between x_{it} and $\bar{x}_{it}$ which can be expressed as:

$$\bar{x}(x_{i,t}) = \phi + \gamma x_{i,t} \tag{18}$$

With $\gamma < 1$, the amenity will be mean-reverting. Therefore, for any two choices of x, denoted x_a and x_b: $|\bar{x}(x_a) - \bar{x}(x_b)| < |x_a - x_b|$. Given this setup, the necessary adjustment to the willingness to pay from the static model is given by:

$$\partial\bar{x}_{i,t}/\partial x_{i,t} = \gamma \tag{19}$$

implying that:

$$\tilde{r}'(\bar{x})\Big|_{\bar{x}=\bar{x}^*_{i,t}} = \frac{1}{\gamma}r'(x)\Big|_{x=x^*_{i,t}} \tag{20}$$

or that the estimate derived by the static model is biased downward by the factor γ. Note that this holds for any level of the trend term, ϕ, which is consistent with the discussion in Section 3.1. Finally, the linearity specified in Equation 18 is done solely for expositional purposes. The adjustment factor (i.e., the derivative given in Equation 19) can easily be allowed to depend upon $x_{i,t}$.

Graphically, the bias can be seen in Figures 6, 7, and 8, where, for illustrative purposes, ϕ is defined so that the mean of x is equal to the mean of $\bar{x}$. In Figure 6a, it can be seen that mean reversion results in a forward-looking housing price function that is everywhere steeper than the one from the static model: $\tilde{r}(\bar{x})$ is increasing at a faster rate than $r(x)$ because

Figure 6: The Amenity is Mean Reverting Through Time

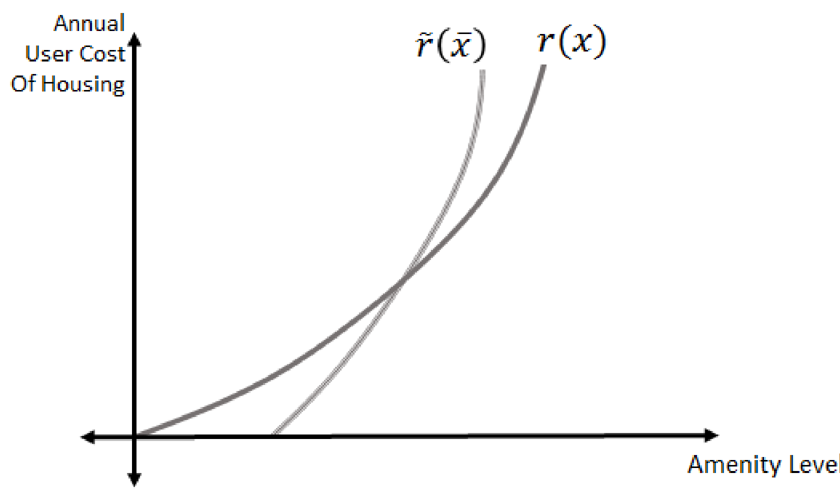

(a) The Housing Price Functions

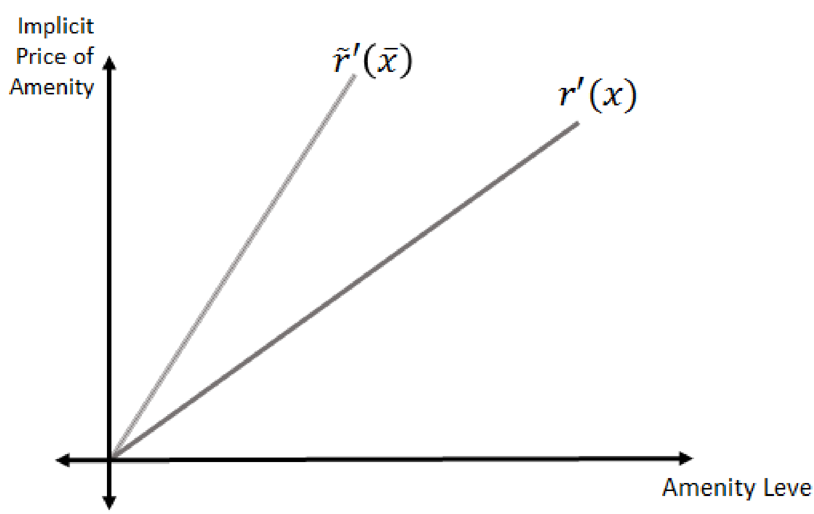

(b) The Implicit Price Functions

each additional dollar spent on housing gets you smaller increase in $\bar{x}$ than it does in x.[15] Correspondingly, $\tilde{r}'(\bar{x})$ lies always above $r'(x)$ as shown in Figure 6b. In other words, the implicit price of the amenity is higher in the forward-looking model.

Figure 7: The Quantity and Price Effects When x is Mean Reverting

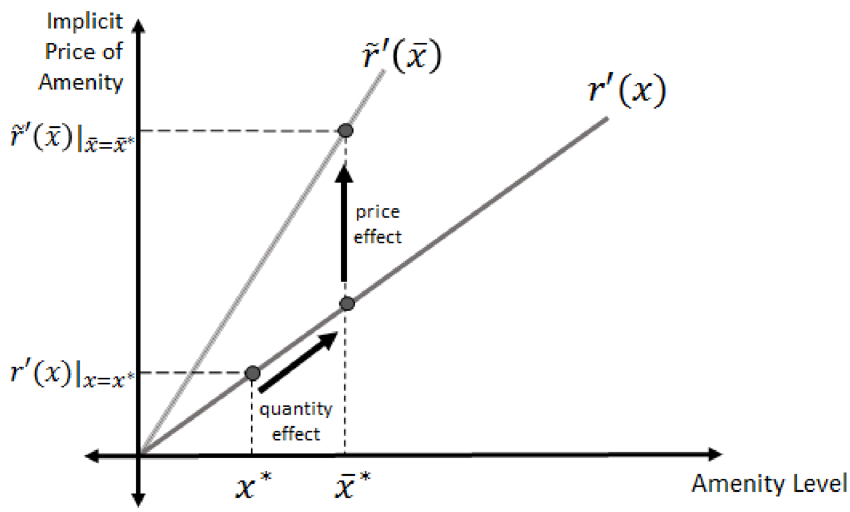

The bias is unambiguously toward zero in the case of mean reversion, regardless of whether or not the overall trend is of increasing or decreasing amenity levels. This is despite the fact that quantity and price effects work in opposite directions for some households in the market. Households that purchase a below-mean level of $x_{i,t}$ will experience a level of $\bar{x}_{i,t}$ that is greater than $x_{i,t}$ and both the quantity and the price effect are positive. This is shown in Figure 7. However, households that purchase an above-mean level of $x_{i,t}$ will experience a level of $\bar{x}_{i,t}$ that is less than $x_{i,t}$; while the price effect is positive, the quantity effect is negative.[16] However, the overall effect is unambiguous; their true willingness-to-pay is higher. This is shown in Figure 8.[17]

[15] In other words, when comparing two houses, the user cost difference is fixed, but the difference reflects a smaller change in $\bar{x}$ than x. So the slope of $\tilde{r}(\bar{x})$ must be larger than $r(x)$ since the horizontal difference is smaller for the same vertical difference.

[16] We have illustrated the case where the mean of x is equal to the mean of $\bar{x}$. More generally, if the amenity is trending through time, mean reversion will imply that low-amenity houses will improve at a faster rate than high-amenity houses.

[17] It can also be seen in Equation 20.

Figure 8: The Quantity and Price Effects When x is Mean Reverting

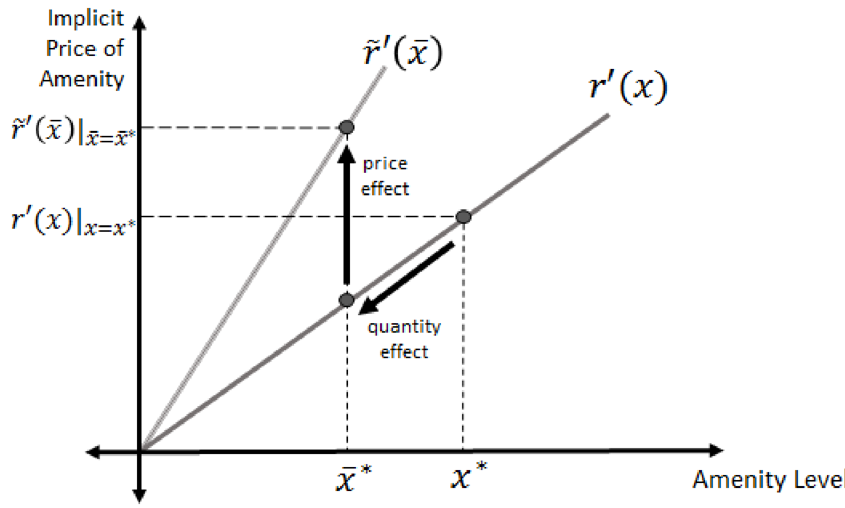

Analogously, with $\gamma > 1$, the amenity will be mean diverging.[18] As before, the mean divergence can take place whether the overall mean of x is rising, falling, or constant. In this case, the willingness to pay derived by the static model will be unambiguously biased away from zero by the factor γ (as $\gamma > 1$). Mean divergence results in a forward-looking housing price function that is everywhere flatter and an implicit price function that is everywhere lower in the forward-looking model.

Households that purchase a below-mean level of $x_{i,t}$ will experience a level of $\bar{x}_{i,t}$ that is less than $x_{i,t}$ and both the quantity and price effects are negative. This can be seen in Figure 9. Households that purchase an above-mean level of $x_{i,t}$ will experience a level of $\bar{x}_{i,t}$ that is greater than $x_{i,t}$; while the price effect is still negative, the quantity effect is positive. However, the overall effect is unambiguous; the static model's estimate of willingness-to-pay is biased away from zero. This is shown in Figure 10.[19]

[18]Therefore, for any two choices of x, denoted x_a and x_b: $|\bar{x}(x_a) - \bar{x}(x_b)| > |x_a - x_b|$.
[19]As in the mean-reversion case, it can also be seen in Equation 20.

Figure 9: The Quantity and Price Effects When x is Diverging

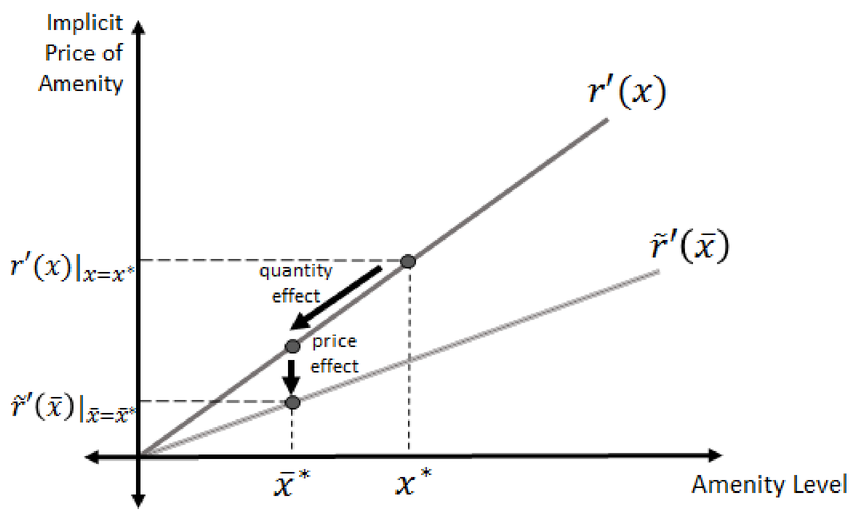

Figure 10: The Quantity and Price Effects When x is Diverging

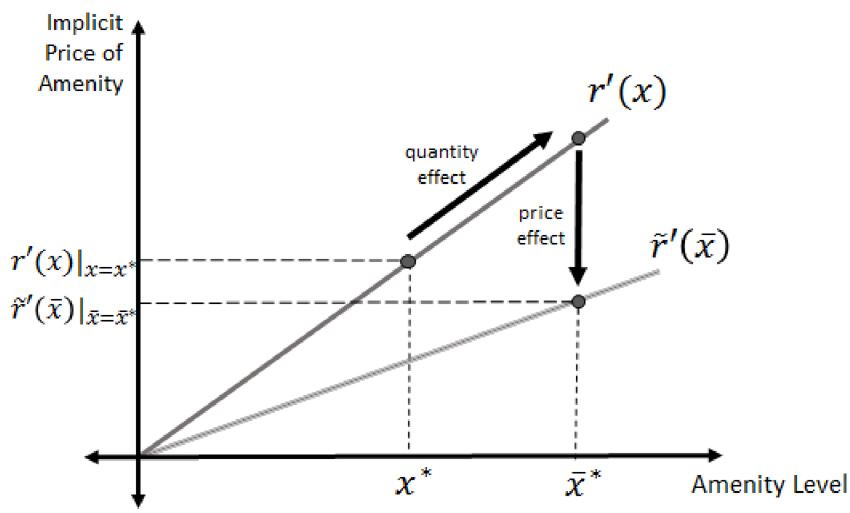

3.3 Nonlinear Utility

In the Appendix, we derive the nonlinear case in greater detail. However, we summarize the key insights here. The effects discussed in Section 3.1 still hold: the price and quantity effects still work in opposite directions. However, due to the nonlinearity of the utility function, they no longer exactly cancel one another out. A natural assumption is that utility is increasing and concave in the amenity. In this case, with an increasing trend in the amenity, the static model will underestimate the marginal willingness to pay. Analogously, with a decreasing trend in the amenity, the static model will overestimate the marginal willingness to pay. Likewise, the effects discussed in Section 3.2 still hold. If the amenity is mean reverting through time, the static model will tend to underestimate the marginal willingness to pay. If the amenity is mean diverging through time, the static model will tend to overestimate the marginal willingness to pay.

4 Application: The Willingness to Pay to Avoid Crime

We now demonstrate the intuition laid out in Sections 2 and 3 in an empirical setting. In particular, we calculate the willingness to pay to avoid violent crime in the Bay Area of California separately by county using the adjustment factor that scales estimates from the static modeling framework.

4.1 Data

In our application, we use a dataset describing housing transactions and violent crime rates for five counties in the Bay Area of California (Alameda, Contra Costa, Marin, San Mateo, and Santa Clara) over the period 1990 to 2008. As this is the same data used in Bishop and Murphy (2011), we can compare the results obtained using the simple adjustment factor approach derived here with the fully dynamic approach used there. Our data are richer than required for illustrating the concepts discussed in this paper, as they allow the econometrician to follow households through time. This richness, however, is needed for the fully-dynamic model.

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, the Census tract in which the house is located, square footage, year built, lot size, and number of rooms.[20] The process of cleaning the data involves a number of cuts which we discuss in more detail in the Appendix.

Crime statistics come from the RAND California database. These data are organized by city and are measured as incidents per 100,000 residents.[21] The data describe annual violent crime rates for each of the cities within the San Francisco Metropolitan area. In this dataset, 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 dataset using an inverse-distance weighted average of the crime rate in each city using the "great circle" calculation. Figure A.1 in the Appendix illustrates the location of these city centroids.

Figures 11 and 12 illustrate the county-specific cross-sectional distributions of violent crime rates and county-specific 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 11: Pooled Cross-Sectional Variation in Violent Crimes per 100,000 Residents

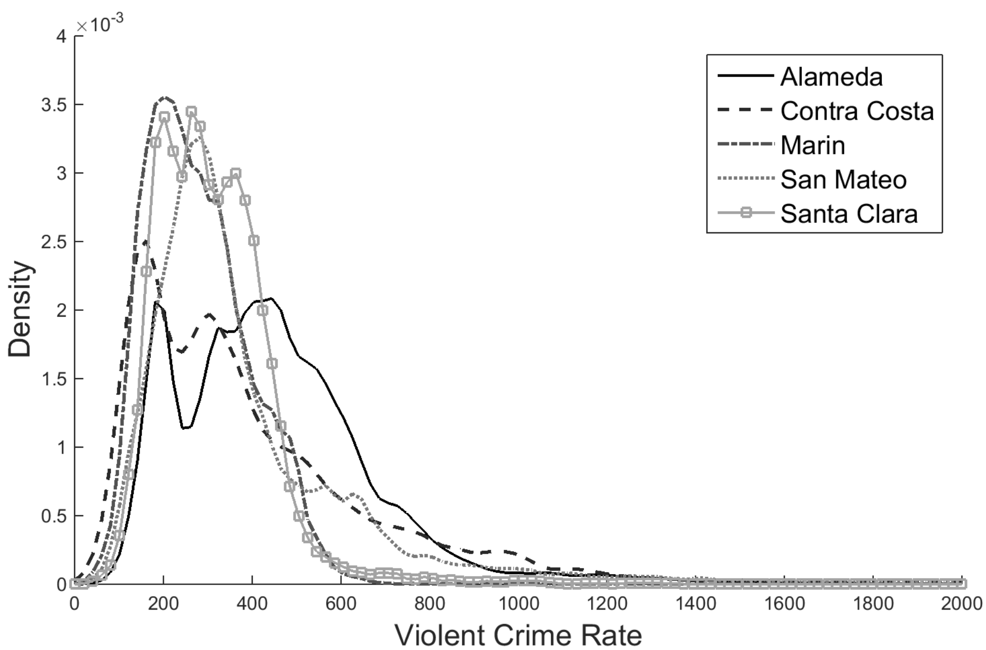

[20]According to the Census Bureau, tracts are small, relatively permanent statistical subdivisions of 1,200 to 8,000 residents.

[21]There are 75 reporting cities in the five counties of analysis.

Figure 12: Pooled Time-Series Variation in Violent Crimes per 100,000 Residents

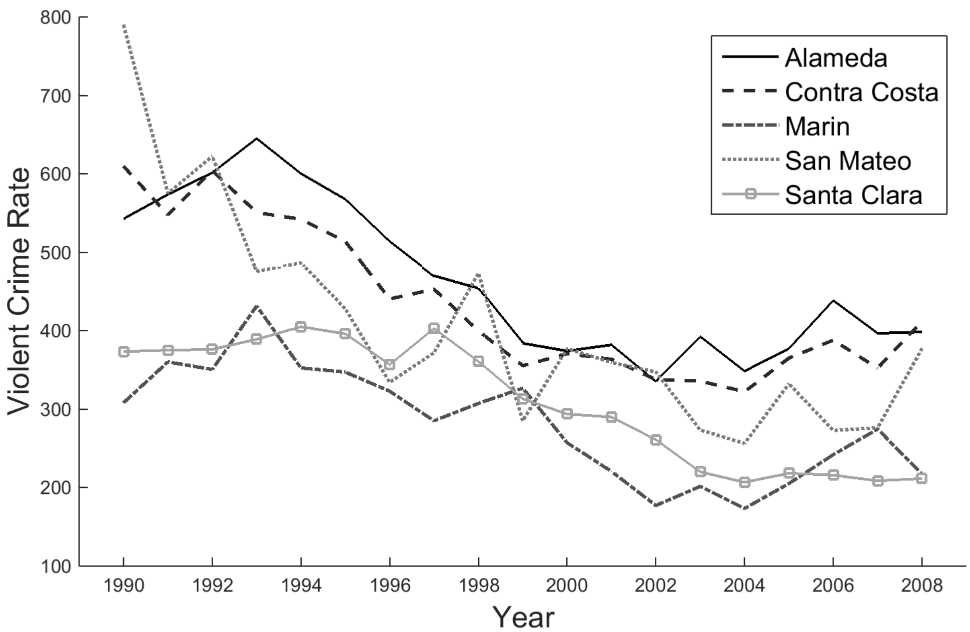

The final sample includes 541,415 transactions which are used to estimate the housing price function separately by county. We then calculate household-specific estimates of marginal willingness to pay, α_i, for each of the 372,334 households in the sample. Summary statistics for both the housing transactions dataset and the merged household dataset may be found in Tables A.1 and A.2 in the Appendix.

4.2 Empirical Specification

While the assumed form of the housing price function is important for correctly estimating households' marginal willingness to pay,[22] the ratio of static and forward-looking estimates is invariant to the choice of functional form and is solely determined by the transition process of the amenity of interest. In this application, we estimate the transition of violent crime by assuming that it follows an AR(1) process:

$$x_{j,k,t} = \rho_{0,k} + \rho_{1,k} x_{j,k,t-1} + \rho_{2,k} t + \varepsilon_{j,k,t} \tag{21}$$

[22]See Cropper, Deck, and McConnell (1988) for a discussion of the choice of functional form in hedonic regressions.

which we estimate separately for each county, k. j denotes the house and t denotes the year of sale. With this simple transition process, $\bar{x}$ may be expressed as $\bar{x}_{j,k,t} = \phi_{k,t} + \gamma_k x_{j,k,t}$ where γ_k is given by the weighted average:

$$\gamma_k = \frac{\sum_{s=1}^{T} \beta^{s-1} \rho_{1,k}^{s-1}}{\sum_{s=1}^{T} \beta^{s-1}} \tag{22}$$

The bias is determined solely by this γ_k parameter following the intuition laid out in Section 3.[23]

We assume the familiar, log specification for the housing price function, which we estimate separately for each county, k:

$$log(r_{j,k,t}) = \theta_{0,k} + \theta_{1,k} x_{j,k,t} + \theta_{2,k} x_{j,k,t}^2 + H'_{j,k,t} \theta_{3,k} + \epsilon_{j,k,t} \tag{23}$$

To convert observed sales prices of houses into annual user costs of housing, we follow the literature and multiply sales prices through by 0.075. The vector of housing attributes, $H_{j,k,t}$, includes property age, square footage, lot size, number of rooms, and a set of dummies for year of sale. We also include a set of fixed effects at the Census-tract level to control for any tract-level, time-invariant unobservables that may be correlated with our measure of violent crime.[24] Taking the exponent of Equation 23 yields $r(x)$, i.e., how the user cost of housing, r, varies with x. As $\bar{x}_{j,k,t} = \phi_{k,t} + \gamma_k x_{j,k,t}$, it is straightforward to recover $\tilde{r}(\bar{x})$, i.e., how the user cost of housing, r, varies with $\bar{x}$, given values of $\phi_{k,t}$ and γ_k.

Each household's marginal willingness to pay to avoid violent crime is then recovered as the value of the implicit price function (i.e., the value of the hedonic gradient) at their chosen level of violent crime exposure. For the static estimates, the annual marginal willingness to pay to avoid one additional crime per $100,000$ residents is recovered according to Equation 5:

$$\hat{\alpha}_i^s = \hat{r}'(x)\Big|_{x=x_i^*}$$

For the forward-looking estimates, the annual marginal willingness to pay to avoid one additional crime per $100,000$ residents is recovered using Equation 15 to adjust each house-

[23]The values of $\phi_{k,t}$ may be recovered analytically, but do not influence the bias.

[24]See Kuminoff, Parmeter, and Pope (2010), who use Monte Carlo evidence to suggest that including spatial fixed effects is the most appropriate way to deal with neighborhood-level unobservables.

hold's static estimate by the county-specific estimate of γ_k:

$$\hat{\alpha}_i^f = \frac{1}{\hat{\gamma}_k}\hat{r}'(x)\Big|_{x=x_i^*}$$

This method is, by construction, equivalent to using Equation 12 directly, following the discussion in Section 3. We set β to 0.95 and T set to seven years.

4.3 Results

We first estimate the transition process for violent crime separately for each of the five counties. The transition probability parameters, i.e., $\rho_{0,k}$, $\rho_{1,k}$, and $\rho_{2,k}$, are reported in Table A.4 in the Appendix. Running the estimation separately by county allows us to calculate the corresponding county-specific values of γ_k according to Equation 22. The estimates of these γ_k parameters, which determine both the size and sign of the bias, are presented in Table 1. It is useful to note that all values of γ_k are strictly less than one, indicating that violent crime is mean reverting through time in each of the counties.[25]

Table 1: County-Specific Estimates of γ

	γ_k
Alameda	0.875
	(.0015)
Contra Costa	0.831
	(.0010)
Marin	0.509
	(.0017)
San Mateo	0.508
	(.0013)
Santa Clara	0.746
	(.0018)

Standard errors in parenthesis

[25]The corresponding values of $\phi_{k,t}$ take the form $\phi_{0,k} + \phi_{1,k}t$ and are shown in Table A.5 in the Appendix. These values imply that, in each of our nineteen years of the sample, expected crime is falling over a seven-year horizon. For Figure 13 (as well as the additional price figures in the Appendix) we use the median value of $\phi_{k,t}$, which is in 1999. However, as discussed in Section 3, this trend term will not affect how our forward-looking model differs from the static model.

We then estimate the housing price function, $r(x)$, separately for each of the five counties. The estimates (and standard errors) are reported in Table A.3 in the Appendix. For exposition, Figure 13 shows both the estimated housing price function and the estimated implicit price function for the most populous county in our sample, Santa Clara.[26] To keep these figures consistent with our earlier theoretical framework (i.e., with first-quadrants plots describing a "good"), we plot how annual user cost of housing varies with the rate of safety, where the safety rate is defined as the negative of the violent crime rate plus a constant (to make the safety rate positive). We plot the housing price functions from the 5th percentile to the 95th percentile of county-specific safety rates, holding the control variables, $H_{j,k,t}$, at their means. Thus, the domains of the plots show both the overall change in safety, as well as the extent of mean reversion in safety rates.[27]

Figure 13: The Housing Price Functions and Implicit Price Functions for Santa Clara

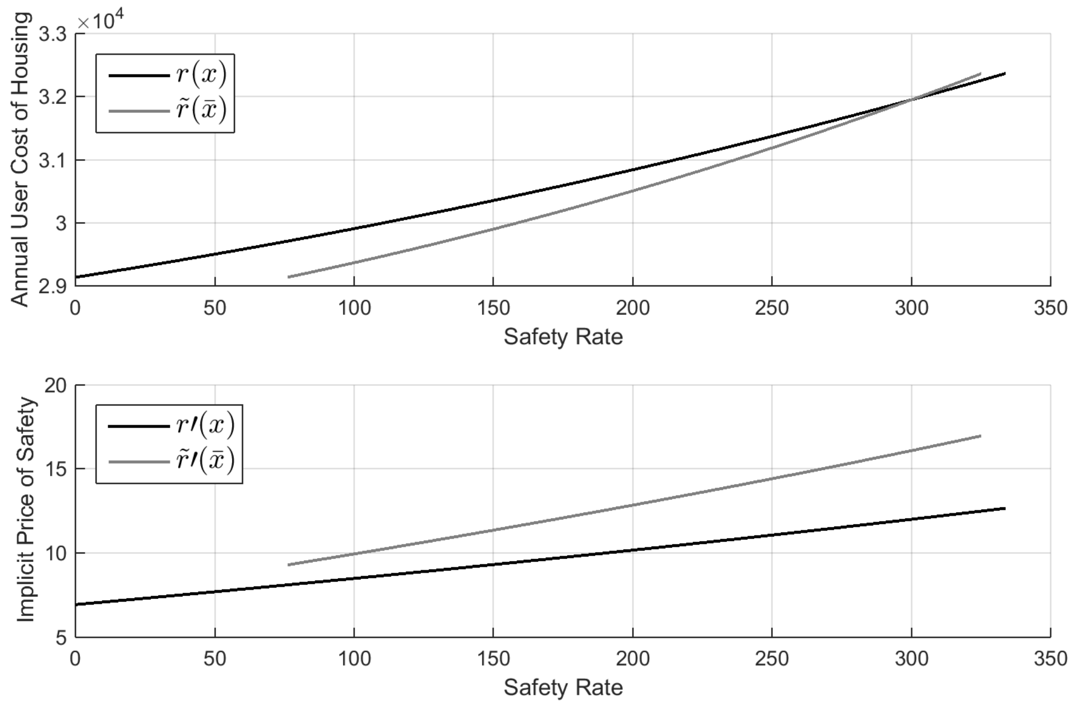

[26]As the price function is estimated very precisely, we don't include confidence intervals in Figure 13. The precision of the estimates may be seen in the small standard errors reported in Table A.3 in the Appendix.

[27]The county-specific constant that we add to make the safety rate positive is the 95th percentile of violent crime. Thus, the lowest value of the safety rate shown in the graphs is, by construction, zero.

It is clear from Figure 13 that the annual user cost of housing is increasing faster in expected average safety than in current safety. In other words, the implicit price of safety is higher than the static model would suggest. Graphically, this can be seen by the fact that $\tilde{r}'(\bar{x})$ lies above $r'(x)$ for all values of x. Figures describing these functions for the other four counties are presented in the Appendix. They appear similar in spirit with the same implied biases and intuition. In all cases, the annual user cost of housing is increasing in safety at an increasing rate (implying an upward-sloping implicit price function).

We find mean reversion in violent crime (and therefore in safety) over the period of our sample. Thus, relative to the mean, households that are currently consuming low levels of violent crime are consuming more on average (in expectation). This translates to a positive quantity effect. Analogously, relative to the mean, households that are currently consuming high levels of violent crime are actually consuming less on average (in expectation). This translates to a negative quantity effect. In both cases, however, households experience a positive price effect; the static model understates the implicit price of violent crime. Thus, the overall effect is that all values of marginal willingness to pay recovered from the static model are biased toward zero.

The static model's estimate of α_i is equal to the estimated implicit price of violent crime (i.e., the value of the hedonic gradient) at each household's observed level of crime exposure. We obtain an estimate for each household in the dataset. The sample mean of the distribution of these estimates is -\$10.66.[28] In other words, the static model implies that the average household dislikes violent crime and is willing to pay \$10.66 per year to avoid one additional crime per 100,000 local residents. This translates to a willingness to pay of \$374.08 per year to reduce total violent crime by ten percent at the mean level of violent crime (350.92 incidents per 100,000 residents).

The forward-looking model's estimate of α_i is obtained by adjusting each household's static estimate by the county-specific term γ_k. As all values of γ_k are estimated to be strictly less than one (i.e., violent crime rates are mean reverting), our forward-looking willingness-to-pay measures will be larger in absolute value. The sample mean of the distribution of these estimates is -\$14.28. The average household is willing to pay \$14.28 per year to avoid one additional crime per 100,000 local residents. This implies that the traditional static approach leads to estimates that are almost thirty-percent lower in absolute terms.

As previously noted, the model laid out here is not fully dynamic in that households may not reoptimize within the seven-year time horizon. However, given the empirical setting,

[28]In less than one percent of cases, $\theta_{1,k} + 2\theta_{2,k}x_{j,k,t} > 0$ which implies positive gradients. We exclude these households from the calculation of utility parameters, which yields utility estimates for 369,015 households.

we are able to directly compare the estimates here to those obtained in Bishop and Murphy (2011). In that paper, a fully dynamic model is estimated using the same data. Interestingly, the forward-looking estimate of willingness to pay found here (-\$14.28) is reasonably close to the fully dynamic estimate of willingness to pay found in Bishop and Murphy (2011) of (-\$13.45).

Table 2: Average Marginal Willingness to Pay by County

	Static MWTP in \$ per year)	Forward-Looking MWTP in \$ per year	Implied Bias in percentage points ($= \gamma_k - 1$)
Alameda	-5.60	-6.40	-12
Contra Costa	-16.54	-19.90	-17
Marin	-16.00	-31.42	-49
San Mateo	-4.24	-8.35	-49
Santa Clara	-11.61	-15.55	-25

Importantly, our specification allows us to explore heterogeneity in the difference in the willingness-to-pay estimates. We estimate all key equations separately for each county in the dataset, recovering county-specific estimates of the parameters of the implicit price function (i.e., the hedonic gradient) and, therefore, of the values of γ_k. Table 2 shows the mean marginal willingness-to-pay estimates separately for each county. The first thing to note is that there are large differences across counties in ratio of static to forward-looking estimates. For example, in Alameda and Contra Costa counties, the static and forward-looking models yield similar results; while crime is falling in both those counties, there is only a small amount of mean reversion, which is the key parameter driving the bias. There is, however, a substantial difference between the estimates of the static and forward-looking models in both Marin and San Mateo counties. While crime is also falling in these counties, there is a considerable amount of mean reversion in crime rates.

Thus, we provide an empirical example which supports the central claims laid out in this paper. As we analyze trends for five separate counties, we find considerable heterogeneity in the size of the bias caused by estimating the static model when the underlying process is forward-looking. That is, (i), even when an amenity is changing over time, the bias generated by specifying the static model may be large or may be small and, (ii), that it is both straightforward and easy to get a sense of this bias through a cursory analysis of the amenity of interest.

5 Conclusion

Researchers in a wide variety of applied fields have relied on Rosen's intuitive 1974 model to recover individuals' marginal willingness to pay for a myriad of implicitly-traded goods and services. In the majority of these applications, researchers have applied the hedonic model to the housing market, recovering estimates of willingness to pay for house- and neighborhood-specific amenities. This housing-market application, however, is also the one most at risk of substantially-biased estimates, given the underlying assumption of free-mobility in the Rosen framework. And, despite many recent advances in the estimation of dynamic models, there continues to exist a substantial burden on the econometrician in terms of both computation and data requirements for the estimation of a dynamic model.

In this paper, we seek to more fully describe the costs and benefits associated with estimating Rosen's familiar model. We illustrate the bias under the assumption that the true, data-generating model is forward-looking using both a series of intuitive graphs and simple algebraic calculations. We then propose a systematic approach to diagnosing the sign and size of the potential bias for a given empirical application based on the time trend of the amenity of interest. We highlight the interesting result where, without reversion to (or divergence from) the average trend through time, there will be no bias (even with changing amenity levels and forward-looking agents). Finally, we propose an adjustment factor which transforms the willingness-to-pay estimate from the static model into that from the forward-looking model.

We highlight these concepts with an empirical application of valuing safety in each of five counties in the Bay Area of California. Employing data describing housing transactions and rates of violent crime, we find considerable heterogeneity (across counties) in the bias associated with specifying the traditional, static model. The results support our suggestion that it may be prudent for the researcher to use a simple analysis of the time-series properties of an amenity to asses the potential benefits prior to adopting a dynamic framework.

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.

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.

Bajari, P., and M. Kahn (2005): "Estimating Housing Demand With An Application to Explaining Racial Segregation in Cities," *Journal of Business and Economic Statistics*, 23(1), 20–33.

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 (2015): "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, forthcoming.

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. (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*, 101(3), 625–629.

Bishop, K. C., and A. Murphy (2015): "Incorporating Dynamic Behavior into the Hedonic Model," 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. (2015): "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.

Cropper, M. L., L. B. Deck, and K. E. McConnell (1988): "On the Choice of Funtional Form for Hedonic Price Functions," *Review of Economics and Statistics*, 70(4), 668–675.

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 (2015): "Neighborhood Choices, Neighborhood Effects, Rents and Moving to Opportunity," 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.

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 tradeoffs 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.

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

Kuminoff, N. V., C. F. Parmeter, and J. C. Pope (2010): "Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities?," *Journal of Environmental Economics and Management*, 60(3), 145–160.

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.

Palmquist, R. B. (1982): "Measuring environmental effects on property values without hedonic regressions," *Journal of Urban Economics*, 11(3), 333–347.

Poterba, J. M. (1984): "Tax Subsidies to Owner-Occupied Housing: An Asset Market Approach," *Quarterly Journal of Economics*, 99(4), 729–52.

Appendix A: Nonlinear Utility

A.1 Static Problem

Households choose x_i to maximize $U(x_i)$, where $U(x_i)$ is given by:

$$U(x_i) = u(x_i) + I_i - r(x_i) \tag{A.1}$$

This yields the first order condition:

$$U'(x_i) = u'(x_i) - r'(x_i) = 0 \tag{A.2}$$

Econometrician can recover $u'(x_i)$ as

$$u'(x)\Big|_{x=x_i^*} = r'(x)\Big|_{x=x_i^*} \tag{A.3}$$

A.2 Forward-Looking Problem

Households choose $x_{i,t}$ to maximize $v(x_{i,t})$, $v(x_{i,t})$ is given by:

$$v(x_{i,t}) = E[\sum_{s=1}^{T} \beta^{s-1}(u(x_{i,t+s-1}) + I_i - r(x_{i,t}))|x_{i,t}] \tag{A.4}$$

we can rewrite (A.4) as:

$$v(x_{i,t}) = \sum_{s=1}^{T} \beta^{t-1} E[u(x_{i,t+s-1})|x_{i,t}] + \sum_{s=1}^{T} \beta^{s-1} I_i - \sum_{s=1}^{T} \beta^{s-1} r(x_{i,t})) \tag{A.5}$$

Choosing $x_{i,t}$ to maximize the value function yields the first order condition:

$$v'(x_{i,t}) = \sum_{s=1}^{T} \beta^{s-1} \frac{\partial E[u(x_{i,t+s-1})|x_{i,t}]}{\partial x_{i,t}} - \sum_{s=1}^{T} \beta^{s-1} r'(x_{i,t}) = 0 \tag{A.6}$$

For notational convenience, let $B = \sum_{s=1}^{T} \beta^{s-1}$. We can then rewrite (A.6) as:

$$r'(x_{i,t}) = \frac{1}{B} \sum_{s=1}^{T} \beta^{s-1} \frac{\partial E[u(x_{i,t+s-1})|x_{i,t}]}{\partial x_{i,t}} \tag{A.7}$$

The bias is determined by whether or not $\frac{1}{B}\sum_{s=1}^{T}\beta^{s-1}\frac{\partial E[u(x_{i,t+s-1})|x_{i,t}]}{\partial x_{i,t}} > \frac{\partial u(x_{i,t})}{\partial x_{i,t}}$. As $\frac{\partial E[u(x_{i,t+s-1})|x_{i,t}]}{\partial x_{i,t}} = \frac{\partial u(x_{i,t})}{\partial x_{i,t}}$ for $s = 1$, we need to compare $\frac{\partial E[u(x_{i,t+s-1})|x_{i,t}]}{\partial x_{i,t}}$ with $\frac{\partial u(x_{i,t})}{\partial x_{i,t}}$ for $s = \{2, 3, ..., T\}$.

The intuition developed in Section 3.2 for mean reverting and mean diverging cases still holds. The conclusion, derived in Section 3.1, that there is no difference between the static and forward looking estimates no longer holds. While the price and quantity effects work still work in opposite directions, they do not exactly offset each other when the utility function is non-linear. We sketch below the direction of the bias for various cases of variance-preserving trends under the assumption that utility is increasing and concave in the amenity.

The first case to consider is where $E[x_{i,t+s-1}|x_{i,t}] = x_{i,t}$. As utility is concave, Jensens's Inequality means $E[u(x_{i,t+s-1})|x_{i,t}] < u(E[x_{i,t+s-1}|x_{i,t}])$. However, as $u' > 0$ and $u'' < 0$, the difference gets smaller as $x_{i,t}$ increases. As such, $E[u(x_{i,t+s-1})|x_{i,t}]$ is steeper in $x_{i,t}$ compared with $u(E[x_{i,t+s-1}|x_{i,t}])$. If $E[x_{i,t+s-1}|x_{i,t}] = x_{i,t}$, then $E[u(x_{i,t+s-1})|x_{i,t}] = u(x_{i,t})$ and $\frac{\partial E[u(x_{i,t+s-1})|x_{i,t}]}{\partial x_{i,t}} > \frac{\partial u(x_{i,t})}{\partial x_{i,t}}$. Therefore, for $E[x_{i,t+s-1}|x_{i,t}] = x_{i,t}$ the static model overstates the marginal willingness to pay.

If $E[x_{i,t+s-1}|x_{i,t}] < x_{i,t}$, there is a decreasing trend in the amenity. In that case, in addition to the effect described in the above paragraph, the concavity of utility will lead $E[u(x_{i,t+s-1})|x_{i,t}]$ to be steeper than $u(x_{i,t})$. This will lead the static model will overestimate the marginal willingness to pay.

Analogously, if there is an increasing trend in the amenity and $E[x_{i,t+s-1}|x_{i,t}] > x_{i,t}$, $E[u(x_{i,t+s-1})|x_{i,t}]$ the concavity of utility will lead $E[u(x_{i,t+s-1})|x_{i,t}]$ to be flatter than $u(x_{i,t})$, which will lead the static model will understate the marginal willingness to pay. If this effect is large enough, it will overcompensate for the $E[x_{i,t+s-1}|x_{i,t}] = x_{i,t}$ effects discussed above and the static model will understate marginal willingness to pay.

Appendix B: Data Cleaning Details

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 assessment, 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 local crime rates using each property's geographic coordinates, we drop properties where latitude and longitude are missing.

A number of additional cuts were made to create the data for Bishop and Murphy (2011). Using the common variables of date, Census tract, 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. We also drop households in the top and bottom 2% based on exposure to crime. 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.

Figure A.1: Cities within the San Francisco Metropolitan Area

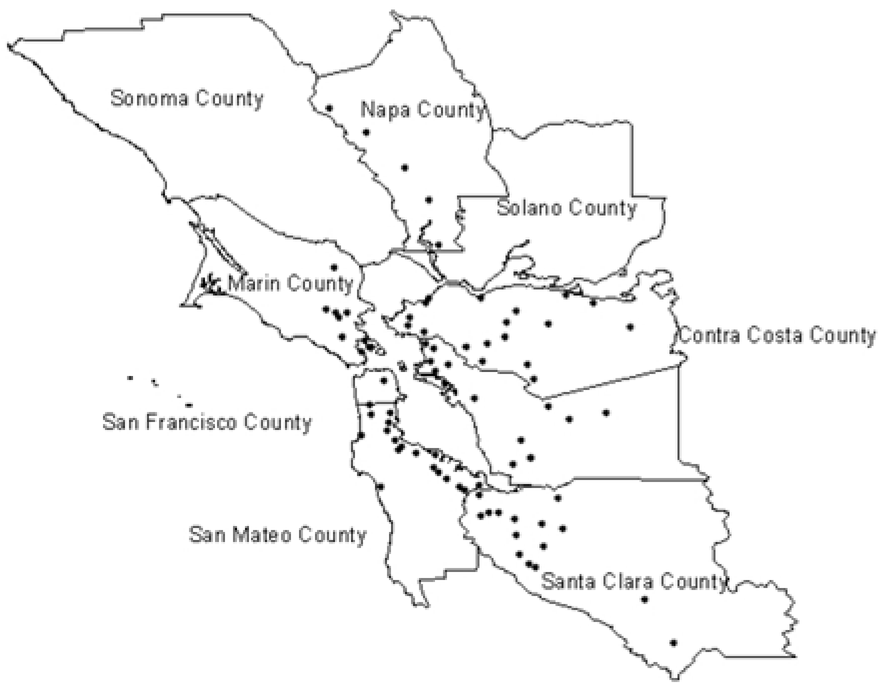

Table A.1: Property Transactions Summary Statistics

N = 541,415					
Variable	Mean	Median	Std. Dev.	Min.	Max.
Sales Price (year 2000 dollars)	434,357	380,138	237,601	75,984	1,662,877
Year of Sale	1999	1999	5.12	1990	2008
Violent Crime Rate (per 100,000 residents)	380	323	263	12.82	3,834
House Square Footage	1,687	1545	662	160	9,130
Lot Square Footage	7,175	6,000	8,034	0	130,680
House Age	32.03	31	20.52	0	147
Number of Rooms	6.68	7	2.33	0	15

Table A.2: Household Summary Statistics

N = 372,334					
Variable	Mean	Median	Std. Dev.	Min.	Max.
Income (year 2000 dollars)	118,825	102,000	68,110	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

Appendix C: Results

Table A.3: Hedonic Regression – Parameter Estimates

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
Violent Crime Rate	-0.20903	-0.70398	-0.96640	-0.12529	-0.46180
	(0.01361)	(0.00803)	(0.08403)	(0.01014)	(0.01977)
Violent Crime Rate Squared	0.02600	0.18787	1.26898	0.02091	0.22928
	(0.00660)	(0.00274)	(0.11632)	(0.00357)	(0.01556)
House Square Footage	0.29287	0.29286	0.37959	0.22190	0.30755
	(0.00140)	(0.00154)	(0.00259)	(0.00230)	(0.00153)
Lot Square Footage	0.00862	0.00826	0.00418	0.00865	0.00654
	(0.00012)	(0.00008)	(0.00018)	(0.00014)	(0.00007)
House Age	-0.00151	-0.00133	0.00090	0.00076	0.00054
	(0.00004)	(0.00005)	(0.00009)	(0.00005)	(0.00003)
Number of Rooms	0.01839	0.03168	0.00143	0.05140	0.03693
	(0.00039)	(0.00050)	(0.00055)	(0.00079)	(0.00048)
Tract Dummies	Yes	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes	Yes
Observations	103,902	138,732	26,255	80,066	192,460
R^2	0.87	0.86	0.75	0.83	0.83

OLS estimates of Equation 23. Standard errors in parentheses.

Violent Crime, House Square Footage, and Lot Square Footage are measured in 1000s of units.

Table A.4: Transition Probability Estimates

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
Constant	30.4772	17.1805	110.6628	94.5984	49.5586
	(0.6115)	(0.4939)	(0.6863)	(1.0404)	(0.4800)
Lagged Violent Crime Rate	0.9518	0.9329	0.7387	0.7375	0.8926
	(0.0006)	(0.0004)	(0.0015)	(0.0011)	(0.0009)
t	-1.2274	0.0473	-3.8823	-0.5867	-2.3091
	(0.0379)	(0.0349)	(0.0334)	(0.0685)	(0.0235)
Observations	3,978,774	3,690,648	611,946	1,762,668	4,551,678
R^2	0.93	0.96	0.69	0.73	0.89

OLS estimates of Equation 21. Standard errors in parentheses.

Table A.5: County- and Year-Specific Estimates of $c_{k,t}$

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
1990	68.49	43.18	178.82	175.18	97.58
	(1.24)	(0.96)	(0.77)	(1.35)	(0.83)
t	-3.23	0.15	-7.20	-1.32	-5.45
	(0.10)	(0.09)	(0.06)	(0.13)	(0.05)

Standard errors in parentheses.

Figure A.2: Alameda

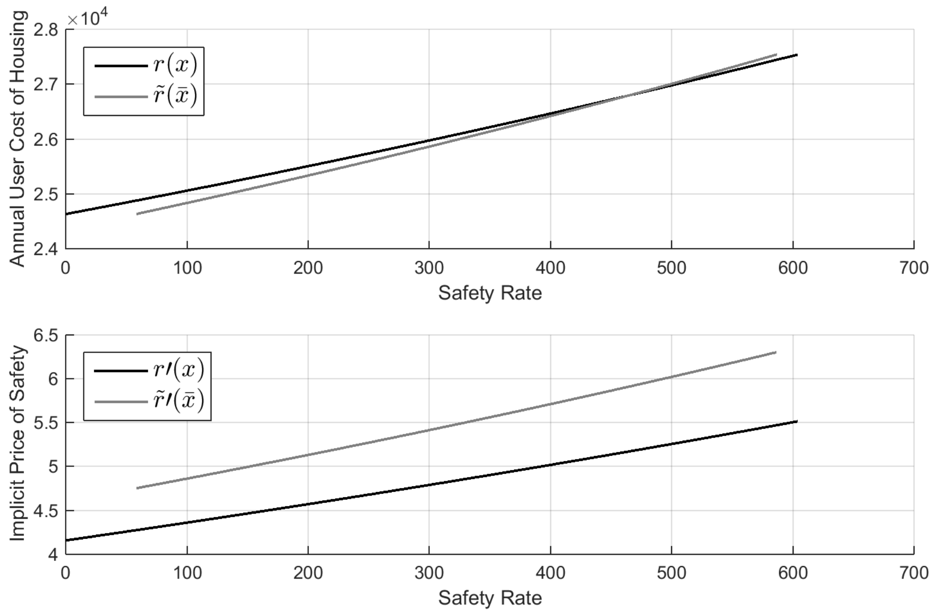

Figure A.3: Contra Costa

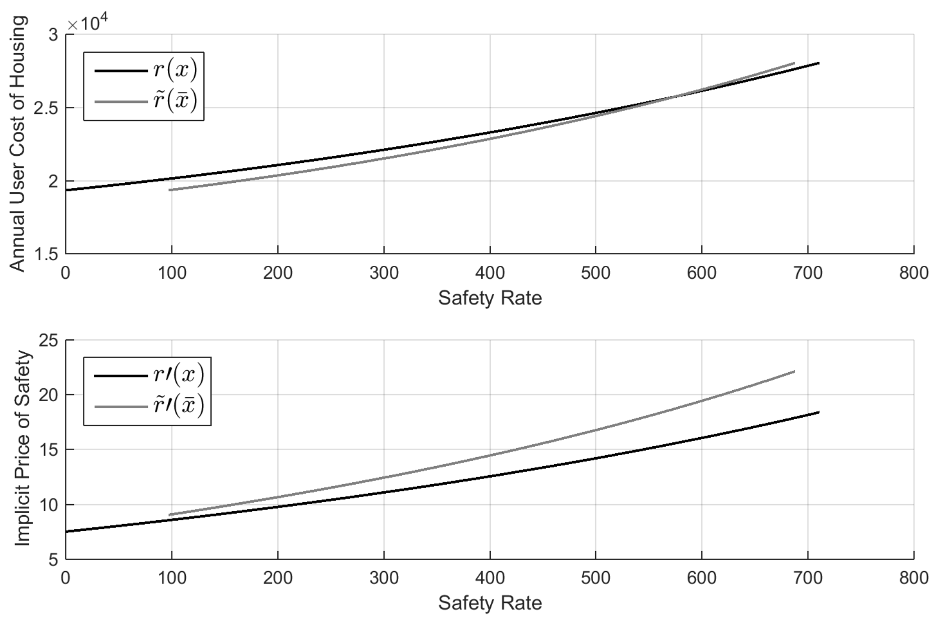

Figure A.4: Marin

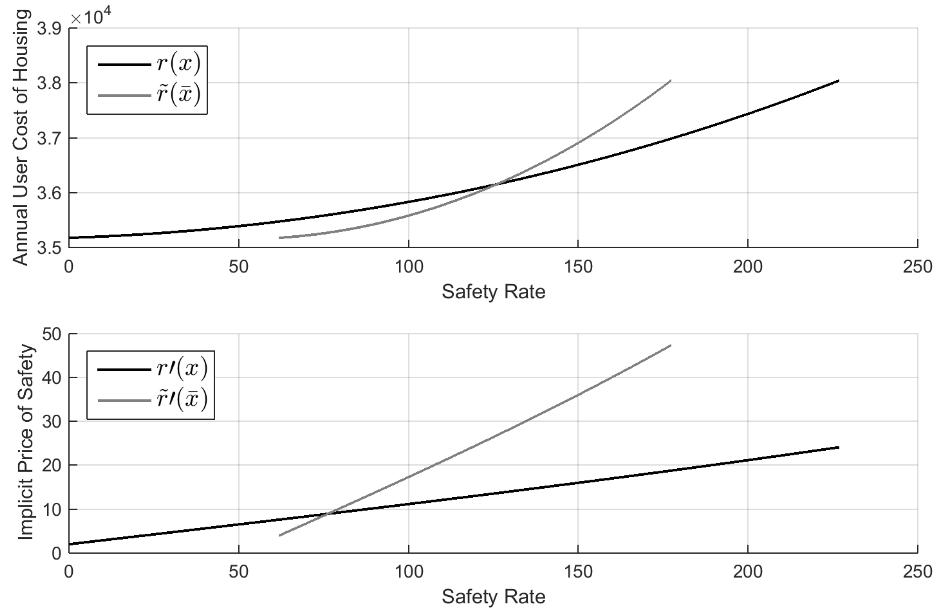

Figure A.5: San Mateo