

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

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 of the amenity of interest. We then derive an adjustment factor that allows the potentially-biased static estimates to be converted into forward-looking estimates. Finally, we motivate the use of this adjustment factor with two empirical applications: the marginal willingness-to-pay to avoid violent crime and the marginal willingness-to-pay to avoid air pollution.

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, Rebecca Diamond, Dennis Epple, Jesse Gregory, Nick Kuminoff, Ralph Mastromonaco, Rob McMillan, Lucija Muehlenbachs, Steve Tadelis, two anonymous referees, as well as various conference and seminar participants for their helpful comments and suggestions. We additionally thank Joshua Murphy for sharing his TSP data.

[†]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, and 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, i.e., 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 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, but not all, cases and the degree of bias will vary substantially.

A recent literature has sought to quantify this bias empirically in specific applications, comparing results from the traditional model with 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 of the amenity of interest. In addition, we illustrate an empirically-relevant example where

¹For recent papers that estimate dynamic models of location choice, see Kennan and Walker (2011), Bishop (2012), Bayer, McMillan, Murphy, and Timmins (2016), Bishop and Murphy (2015), Caetano (2016), Davis, Gregeory, Hartley, and Tan (2017), and Mastromonaco (2015).

the static model and the forward-looking model arrive at the same estimate of willingness to pay, despite the fact that the amenity of interest is time-varying.

To understand the potential bias arising from a static model, it is worthwhile revisiting the intuitive identification strategy of the static Rosen model that allows the researcher to recover estimates of marginal willingness to pay for an amenity using information on (i) the currently-observed quantity of the amenity that the household chooses to consume and (ii) the price schedule of the amenity faced by the household. However, if households are instead choosing a house based on some average stream of future amenities (and not solely on the currently-observed levels), the traditional model will get both (i) and (ii) wrong, resulting in potentially biased estimates of marginal willingness to pay. It is straightforward to see that by using the incorrect interpretation of quantity, the static model will either under- or over-attribute the true quantity “consumed”. 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 schedule from housing price differentials using the incorrect interpretation of quantity, 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. Additionally, based on these time-series properties, we show that an adjustment factor may be easily derived for a given empirical application. This adjustment factor may be used to convert the estimates of marginal willingness to pay from a static model into those that one would have obtained using our forward-looking model. In the simplest case, this adjustment factor is a constant and can be recovered from a simple, ordinary least squares regression.

In the first of our empirical applications, we use a rich data set on housing transactions to apply our adjustment factor to illustrate an example with large heterogeneity in the size of the bias associated with the static model. In particular, we use data from the San Francisco Bay Area to estimate the marginal willingness to pay to avoid violent crime. We calculate the adjustment factor 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. 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 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.

In our second application, we illustrate how our adjustment factor approach may easily

be applied to existing estimates of the hedonic price function. In particular, we transform the housing-market estimates from Chay and Greenstone’s well-known, 2005 paper to recover static and forward-looking measures of marginal willingness to pay to avoid air pollution in the years following the Clean Air Act Amendments of 1970. This application, which focuses on a period with a policy-driven, mean-reverting decline in air pollution, also serves as an example of a case with a large adjustment factor.

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 static 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 in the two empirical settings; and Section 5 concludes.

2 Model

In this section, we provide an overview of the traditional, static model of Rosen (1974), as well as a simple, forward-looking model of marginal willingness to pay.

2.1 The Traditional, Static Model of Willingness to Pay

We first consider the static model of marginal willingness to pay for a house or neighborhood amenity. In this model, households maximize current utility with respect to their choice of amenity consumption. 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 .² Broadly speaking, the intuition developed here applies to non-linear specifications, which we present in the Appendix.

²While 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.

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 user cost of housing as capturing either a rent or mortgage payment.³ The function that relates the level of amenity consumption, x , to the annual user cost of housing, r_i , is the housing price function, $r_i = r(x_i)$. We do not need to make any assumptions regarding the functional form of the housing price function, $r(x)$.⁴

Incorporating the household's budget constraint, i.e., that their numeraire consumption, C_i , is equal to income, I_i , minus the 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 .⁵ 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 the amenity x . 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.⁶ This naturally suggests the (static) estimator:

$$\hat{\alpha}_i^s = \hat{r}'(x) \Big|_{x=x_i^*}, \tag{4}$$

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

³The 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.

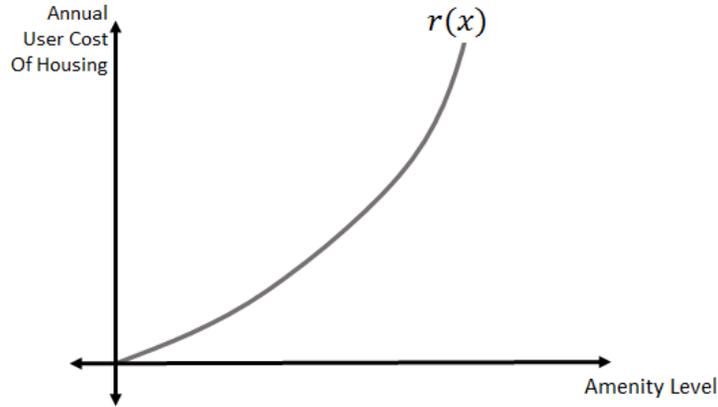
⁴As we do not specify the preferences of housing suppliers, the linear-utility assumption could be consistent with any equilibrium price function that is strictly convex. See Bajari and Benkard (2005) for a discussion. Furthermore, the results derived here hold for any such price function. When we draw the price function, we must show a particular form and, for simplicity, we illustrate a quadratic housing price function (linear implicit price function) in Figures 1, 2, and 3.

⁵If x were 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.

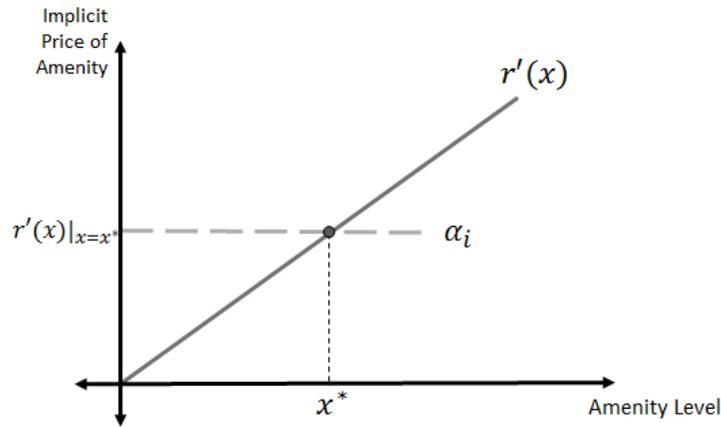
⁶The approach is simple as it only requires interpreting the estimated implicit price function as the menu of prices households face. Given this menu of prices, each household's amenity choice reveals its preferences, which are recovered separately for each household in the data. The approach is limited, however, as without modeling the preferences of housing suppliers, we cannot speak to the equilibrium-price effects of counterfactual amenity changes, as in Cellini, Ferreira, and Rothstein (2010).

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

Figure 1: Graphical Representation of the Model



(a) The Housing Price Function



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

Using each family's observed consumption of x and the implicit price function, the econometrician is able to construct the implicit price of x that each family actually paid. This information, paired with the first-order condition for utility maximization, allows the econometrician to invert the implicit price and recover the household-specific preference

parameter, α_i . This inversion is depicted in Figure 1(b).⁷

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.⁸

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 an 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. 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, i.e., 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.⁹

⁷Note that the recovery of the implicit price function and mapping into a local measure of willingness to pay for each household is commonly referred to as the first stage of Rosen's 1974, two-stage model. Applying the assumption that each household has a flat marginal willingness-to-pay function, this inversion would also serve as the second stage of Rosen's model: the recovery of a global measure of willingness to pay. See Bajari and Benkard (2005) and Bajari and Kahn (2005) for an insightful discussion of the interpretations of α_i as either a structural parameter of the utility function or as a local estimate of marginal willingness to pay at the point of consumption.

⁸If agents are actually myopic (for whatever reason), we assume that the static framework presented here yields the correct estimate of α_i . However, an alternative view of this modeling framework might be that x serves as a proxy for the associated future stream of amenities and the recovered $\hat{\alpha}_i$ may be interpreted as the per-period marginal willingness to pay when households are forward-looking. In this case, x would suffer from measurement error and the insights in Chalfin and McCrary (2015), for example, would apply.

⁹In 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.

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 the current choice of 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.¹⁰

The amenity of interest, x , is evolving through time and households form expectations over future levels of x . We write the discounted sum of annual utility flows over the next T years (i.e., the value function) as:

$$v(x_{i,t}) = E\left[\sum_{s=1}^T \beta^{s-1} (\alpha_i x_{i,t+s-1} + I_i - r(x_{i,t})) \mid x_{i,t}\right] \quad (5)$$

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

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

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, $r_{i,t}$:

$$r_{i,t} = \tilde{r}(\bar{x}(x_{i,t})) = r(x_{i,t}) \quad (6)$$

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

Defining $\tilde{v}(\bar{x})$ analogously (i.e., $\tilde{v}(\bar{x}(x_{i,t})) = v(x_{i,t})$), allows us to rewrite Equation (5) in terms of $\bar{x}_{i,t}$:

$$\tilde{v}(\bar{x}_{i,t}) = \alpha_i \sum_{s=1}^T \beta^{s-1} \bar{x}_{i,t} + \sum_{s=1}^T \beta^{s-1} I_i - \sum_{s=1}^T \beta^{s-1} \tilde{r}(\bar{x}_{i,t}) \quad (7)$$

and the household's problem is then equivalent to choosing $\bar{x}_{i,t}$ to maximize $\tilde{v}(\bar{x}_{i,t})$, yielding

¹⁰A 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 first-order condition:¹¹

$$\tilde{v}'(\bar{x}_{i,t}) = \alpha_i \sum_{s=1}^T \beta^{s-1} - \sum_{s=1}^T \beta^{s-1} \tilde{r}'(\bar{x}_{i,t}) = 0 \quad (8)$$

The first-order condition described by Equation (8) may then be used to solve for household i 's marginal willingness to pay for amenity x . This naturally suggests the (forward-looking) estimator:

$$\hat{\alpha}_i^f = \hat{r}'(\bar{x}) \Big|_{\bar{x}=\bar{x}_{i,t}^*} \quad (9)$$

for the per-annum marginal 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 (4)), Equation (9) highlights the two effects that we previously referred to as the price effect and the quantity effect.¹² 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 Figure 1, 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 reoptimization is not without cost, estimates of marginal 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 (4)) and the estimate of marginal willingness to pay from the forward-looking model (Equation (9)).

¹¹If 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 (8).

¹²This highlights the fact that our framework focuses on the mapping of implicit prices into preference parameters; if households are forward looking, using Equation (4) to recover α_i (instead of Equation (9)) would be a case of model mis-specification. In other words, we abstract from the issues associated with inconsistent estimates of implicit prices (which would affect either modeling framework). For a discussion of these issues, see Chay and Greenstone (2005) and Bajari, Fruehwirth, Kim, and Timmins (2012) (both with hedonics-related applications) and Kane and Staiger (2008) (more generally).

The transition properties of the amenity of interest will determine both 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 Simply 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} may be expressed:

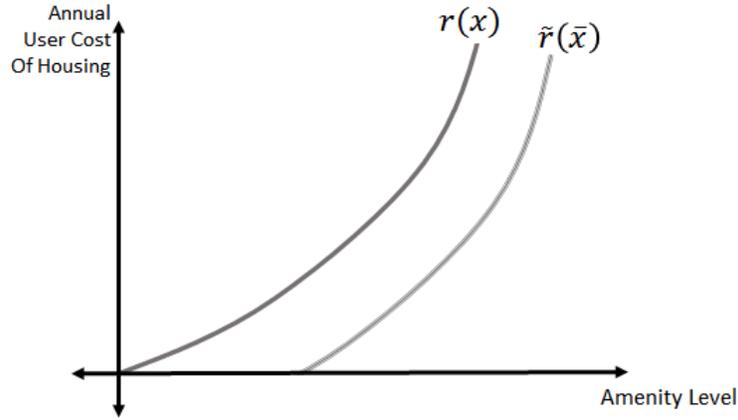
$$\bar{x}(x_{i,t}) = \phi + x_{i,t} \tag{10}$$

where ϕ can be either positive or negative and with 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 relationship ensures that a change in current $x_{i,t}$ produces a one-for-one change in average future amenity consumption, \bar{x} , i.e., for any two choices of current x , denoted x_a and x_b : $|\bar{x}(x_a) - \bar{x}(x_b)| = |x_a - x_b|$.

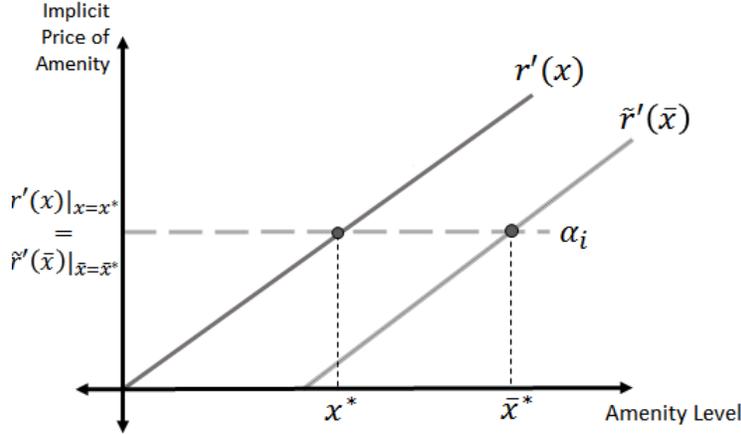
Graphically, a uniform increase in x ($\phi > 0$) is represented in Figure 2. In Figure 2(a), one can see that the increase in x results in a forward-looking price function that lies (in a parallel manner) to the right of the static price function. In other words, for any given level of housing expenditure, 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 2(b), lies (in a parallel manner) to the right of the static implicit price function. In other words, the implicit price of the amenity is lower than the static model would imply.

As can be noted in Figure 2(b), the quantity effect (using \bar{x} instead of x) and the price effect (using $\tilde{r}'(\cdot)$ instead of $r'(\cdot)$) work in opposite directions and *exactly* offset one another; there is no bias associated with the static modeling framework, even for an amenity with

Figure 2: The Amenity is Rising Through Time



(a) The Housing Price Functions



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

potentially large time-trends.

This holds true for the analogous case of a uniform decrease in x , i.e., $\phi < 0$. In this case, the decrease in x will result in a forward-looking price function that lies (in a parallel manner) to the left of the static price function and an implicit price function that lies (in a parallel manner) to the left of the static implicit price function. The negative quantity effect will be perfectly offset by the positive price effect, and there will be no bias.

In addition to a graphical representation, we can analytically derive an expression for the difference between the static and forward-looking estimates and show that this adjustment factor will equal zero when the amenity is simply rising or falling through time. Using the definition that appears in Equation (6), i.e., that $\tilde{r}(\bar{x}(x_{i,t})) = r(x_{i,t})$, and differentiating with

respect to $x_{i,t}$ yields:

$$\tilde{r}'(\bar{x}_{i,t}) \frac{\partial \bar{x}_{i,t}}{\partial x_{i,t}} = r'(x_{i,t}) \quad (11)$$

implying that:

$$\underbrace{\tilde{r}'(\bar{x}) \Big|_{\bar{x}=\bar{x}_{i,t}^*}}_{\text{MWTP forward-looking model}} = \frac{1}{\frac{\partial \bar{x}_{i,t}}{\partial x_{i,t}}} \underbrace{r'(x) \Big|_{x=x_{i,t}^*}}_{\text{MWTP static model}} \quad (12)$$

First, it can be easily seen in Equation (12) that a simple adjustment factor captures the difference between the marginal willingness-to-pay estimate from the static model and that from the forward-looking model. Second, it can be seen that when $\bar{x}(x_{i,t}) = \phi + x_{i,t}$, $\bar{x}_{i,t}/\partial x_{i,t} = 1$ and the marginal willingness to pay derived using the static model will be identical to that derived using the forward-looking one.

The derivative, $\partial \bar{x}_{i,t}/\partial x_{i,t}$ presents a simple and practical adjustment to estimates from the static model, i.e., one could recover this adjustment factor in a separate first-stage and divide $\hat{\alpha}_i^s$ by $\partial \bar{x}_{i,t}/\partial x_{i,t}$. This adjustment-factor approach will be used in our empirical applications to convert estimates from the static model of marginal willingness to pay into ones that would have obtained using the forward-looking model.

Finally, we note that this analysis also applies to the multivariate case where x is simply treated as a vector of amenities.¹³

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 would 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. An example of mean diversion would be tipping points; for example, in the case of neighborhood racial composition, neighborhoods with minority levels above some tipping point may experience more in-migration of minorities. For our purposes, however, we do not distinguish between the underlying causes of mean reversion or mean divergence.

¹³Some minor notational adjustments would be required. Let α_i , $r'(x_{i,t})$, and $\tilde{r}'(\bar{x}_{i,t})$ be row vectors and x be a column vector. The adjustment factor, $\frac{\partial \bar{x}_{i,t}}{\partial x_{i,t}}$, would be a (Jacobian) matrix and Equation (12) would read: $\tilde{r}'(\bar{x}) \Big|_{\bar{x}=\bar{x}_{i,t}^*} = r'(x) \Big|_{x=x_{i,t}^*} \left[\frac{\partial \bar{x}_{i,t}}{\partial x_{i,t}} \right]^{-1}$.

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

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

With $\gamma < 1$, the amenity will be mean-reverting, i.e., for any two choices of x , denoted x_a and x_b : $|\bar{x}(x_a) - \bar{x}(x_b)| < |x_a - x_b|$. With $\gamma > 1$, the amenity will be mean-diverging, i.e., $|\bar{x}(x_a) - \bar{x}(x_b)| > |x_a - x_b|$.

Given this setup, the necessary adjustment factor to the willingness to pay from the static model is given by:

$$\partial \bar{x}_{i,t} / \partial x_{i,t} = \gamma \quad (14)$$

implying that:

$$\tilde{r}'(\bar{x}) \Big|_{\bar{x}=\bar{x}_{i,t}^*} = \frac{1}{\gamma} r'(x) \Big|_{x=x_{i,t}^*} \quad (15)$$

or that the estimate derived by the static model is biased 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 (13) is done solely for expositional purposes; the adjustment factor (i.e., γ in Equation (14)) can easily be allowed to depend on $x_{i,t}$.

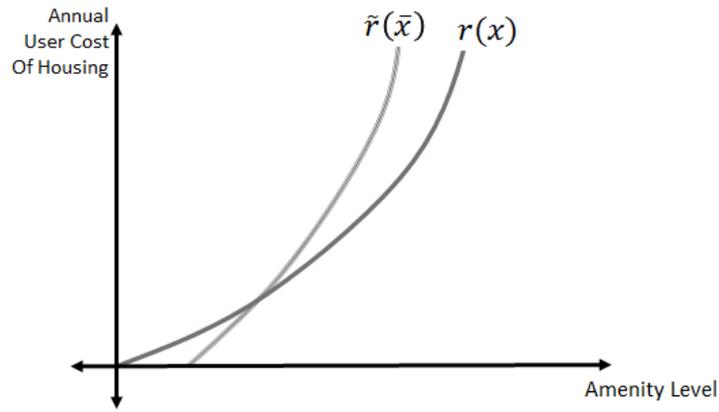
In Figure 3(a), we show the price functions for mean-reversion of the amenity. It may be seen that the forward-looking housing price function is everywhere steeper than the one from the static model; $\tilde{r}(\bar{x})$ is increasing at a faster rate than $r(x)$, as each additional dollar spent on housing buys the household a smaller increase in \bar{x} than it does in x .¹⁴ Correspondingly, $\tilde{r}'(\bar{x})$ lies always above $r'(x)$, as shown in Figures 3(b) and 3(c). In other words, the implicit price of the amenity is higher in the forward-looking model.

The bias is unambiguously toward zero in the case of mean reversion. 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 effects will be positive. This is shown in Figure 3(b). Households that purchase an above-mean level of $x_{i,t}$ will experience a level of $\bar{x}_{i,t}$ that is lower than $x_{i,t}$; while the price effect will be positive, the quantity effect will be negative. This is shown in Figure 3(c).¹⁵ However, the overall effect is unambiguous; as be seen analytically in Equation (15), their true willingness-to-pay is higher.

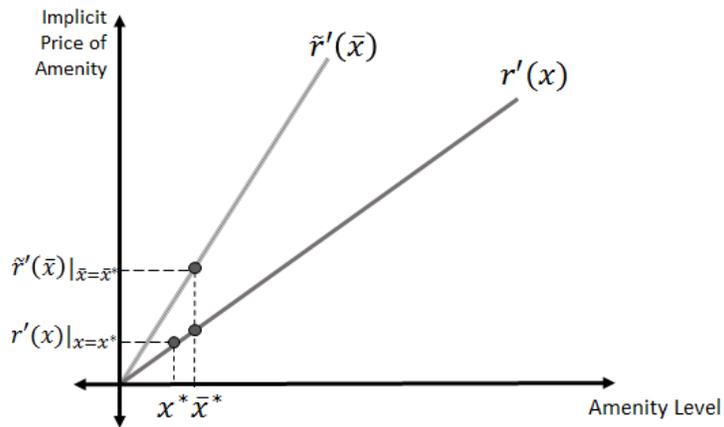
¹⁴In other words, when comparing two houses, the user-cost difference is fixed, but the user-cost difference reflects a smaller change in \bar{x} than x . Thus, the slope of $\tilde{r}(\bar{x})$ must be larger than $r(x)$ as the horizontal difference is smaller for the same vertical difference.

¹⁵In Figures 3(b) and 3(c), 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.

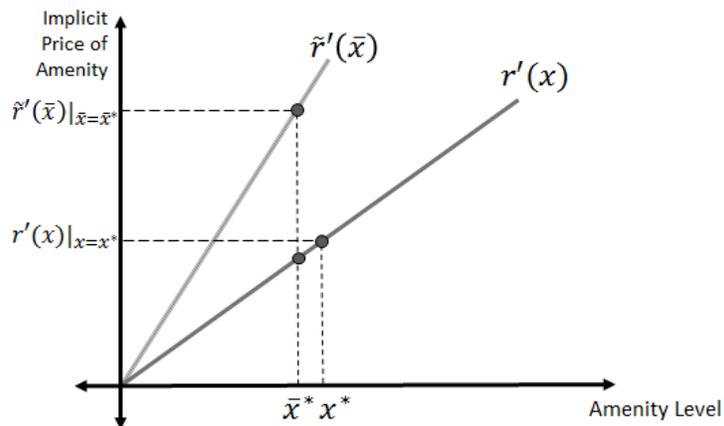
Figure 3: The Amenity is Mean Reverting Through Time



(a) The Housing Price Functions



(b) The Implicit Price Functions and Household i 's MWTP Below-Mean Choice of x^*



(c) The Implicit Price Functions and Household i 's MWTP Above-Mean Choice of x^*

Analogously, when the amenity is mean diverging, the willingness to pay derived by the static model will be unambiguously biased away from zero by the adjustment factor γ (as $\gamma > 1$). Mean divergence results in a forward-looking price function that is everywhere flatter and a forward-looking implicit price function that is everywhere lower than in the static model.

3.3 Nonlinear Utility

In the Appendix, we derive the nonlinear case in greater detail, including an empirical specification where utility is a function of $\log(x)$.¹⁶ The two key insights are as follows. First, the effects of a simply rising/falling amenity discussed in Section 3.1 still hold, broadly speaking: the price and quantity effects still work in opposite directions, yet they no longer *exactly* cancel one another out, due to the nonlinearity of the utility function. (With concave utility, the static model will underestimate (overestimate) willingness to pay when the amenity is increasing (decreasing) over time.) Second, the effects of a mean-reverting/mean-diverging amenity discussed in Section 3.2, still hold, broadly speaking: the static model will underestimate (overestimate) willingness to pay when the amenity is mean-reverting (mean-diverging) over time.

4 Empirical Applications

We now demonstrate the intuition laid out in Sections 2 and 3 in two empirical settings. In the first application, of recovering the marginal willingness to pay to avoid violent crime, we highlight a setting with large heterogeneity in the adjustment factor. In the second application, of recovering the willingness to pay to avoid air pollution, we revisit the analysis of a major U.S. policy: the 1970 Clean Air Act Amendments.

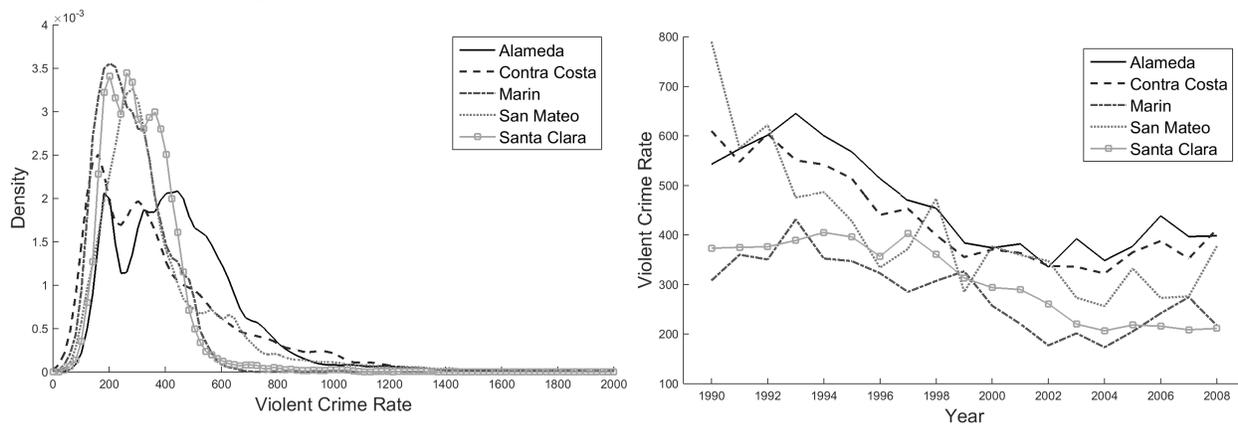
4.1 The Willingness to Pay to Avoid Crime

For this application, we begin with a dataset describing housing transactions and violent crime rates for five counties located in the Bay Area of California (Alameda, Contra Costa, Marin, San Mateo, and Santa Clara) over the period 1990 to 2008. The housing transactions data were purchased from DataQuick and contain dates of sale, prices, geographical coordinates, property age, square footage, lot size, and number of rooms (in addition to a number

¹⁶A well-known limitation of the linear-utility model is that the degree of uncertainty over future amenity values does not affect the current optimal choice and, as such, risk aversion plays no role in decision making.

of other variables). These transaction data are much richer than required for illustrating the concepts discussed in this paper, as they allow the econometrician to follow households through time (note that this richness would be needed for the fully-dynamic model). The crime statistics come from the FBI Uniform Crime Reporting Program and were accessed via the RAND California database. These data are organized by city and are measured as incidents per 100,000 residents. The data describe annual violent crime rates for seventy-five cities within the San Francisco Metropolitan area.¹⁷ 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 city-level crime rates using the “great circle” calculation.¹⁸ 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 for each of the 372,334 households in the sample.¹⁹ Respectively, the panels in Figure 4 illustrate the county-specific, distributions and time-series of the violent crime rate.

Figure 4: Variation in Violent Crimes per 100,000 Residents



Employing these data to recover the parameters of the housing price function, we adopt the familiar, log-linear specification, which we estimate separately for each county, k :²⁰

$$\log(r_{j,k,t}) = \theta_{0,k} + \theta_{1,k}x_{j,k,t} + \theta_{2,k}x_{j,k,t}^2 + H'_{j,k,t}\theta_{3,k} + \epsilon_{j,k,t} \quad (16)$$

¹⁷Figure A.1 in the Appendix illustrates the location of these city centroids.

¹⁸This algorithm finds the shortest distance between any two points on the surface of a sphere.

¹⁹The process of cleaning these data involves a number of cuts. In the Appendix, we discuss these cuts and present summary statistics for both the housing transactions dataset and the household dataset in Tables A.2 and A.3, respectively.

²⁰While the choice of functional form for this price function is important for correctly recovering households’ marginal willingness to pay, the ratio of static willingness to pay and forward-looking willingness to pay estimates is invariant to this choice, as discussed in Section 3. For a complete discussion of the choice of functional form in hedonic regressions, see Cropper, Deck, and McConnell (1988).

We follow the literature and multiply observed sales prices by 0.075 to convert to annual user-costs of housing. The vector of housing attributes, $H_{j,k,t}$, includes property age, square footage, lot size, number of rooms, a set of dummies for year of sale, and a set of Census-tract fixed effects to control for any tract-level, time-invariant unobservables that may be correlated with our measure of violent crime.²¹ The results are as expected: housing price is decreasing at a decreasing rate in violent crimes for all five counties.²²

As laid out in Section 3, the ratio of static and forward-looking estimates of marginal willingness to pay is solely determined by the transition process of the amenity of interest. In this application, we assume that households have rational expectations and specify the transition of violent crime as following an AR(1) process. We estimate the following equation:²³

$$x_{j,k,t} = \rho_{0,k} + \rho_{1,k}x_{j,k,t-1} + \rho_{2,k}t + \varepsilon_{j,k,t} \quad (17)$$

where k denotes county, j denotes the house, and t denotes the year of sale. With this simple transition process, \bar{x} may then 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:

$$\gamma_k = \frac{\sum_{s=1}^T \beta^{s-1} \rho_{1,k}^{s-1}}{\sum_{s=1}^T \beta^{s-1}} \quad (18)$$

where we set β to the commonly-used discount factor of 0.95 and set T to seven years, which is approximately the median household tenure in the United States over our sample period. Our results are robust to each of these choices, with the obvious caveat that a very large decrease in either T or β would violate the intuition laid out in Sections 2 and 3, as either $T = 1$ or $\beta = 0$ collapse the model into the static framework.²⁴ To calculate standard errors, we implemented a bootstrap procedure with 1000 draws. Standard errors for $\hat{\gamma}$ could also be obtained using the delta method.²⁵

Calculated according to Equation (18), the county-specific estimates of the adjustment

²¹According to the Census Bureau, Census tracts are small, relatively permanent statistical subdivisions of 1,200 to 8,000 residents. 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 omitted-variable bias from neighborhood-level unobservables.

²²The price-function estimates (and standard errors) for each of the five counties are reported in Table A.6 in the Appendix. These price functions (and associated implicit price functions) are also presented in Figure A.2 in the Appendix.

²³Our results are effectively unchanged when we consider alternative specifications for the formation of expectations. Results from both an adaptive-expectations framework and an AR(2) framework, as well as a specification with an additional control amenity, are presented in Table A.7 in the Appendix.

²⁴Results for $T \in \{5, 9\}$ and $\beta \in \{0.90, 1.00\}$ are presented in Table A.7 in the Appendix.

²⁵Using the delta method, the standard errors of $\hat{\gamma}$ can be constructed as $\partial\gamma/\partial\rho_1$ times the standard error of ρ_1 , where $\partial\gamma/\partial\rho_1 = (\sum_{s=1}^T \beta^{s-1})^{-1} \sum_{s=1}^T \beta^{s-1}(s-1)\rho_1^{s-2}$.

factors, γ_k , which determine both the size and sign of the bias, are presented in Table 1.²⁶ All values of γ_k are strictly less than one, indicating that violent crime is mean-reverting in each of the five counties.²⁷

Table 1: County-Specific Estimates of γ

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
γ_k	0.8751	0.8313	0.5091	0.5077	0.7465
	(0.0003)	(0.0002)	(0.0009)	(0.0005)	(0.0004)

Standard errors are given in parentheses.

Using the estimates from the price function alone, we recover the static model’s estimate of annual marginal willingness to pay according to Equation (4), i.e., $\hat{\alpha}_i^s$ is equal to the estimated implicit price of violent crime at each household’s observed level of violent crime exposure.²⁸ The sample mean of the distribution of these estimates is -\$10.66 (in year-2000 dollars) with a standard error of \$0.19. 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 an average willingness to pay of \$37.41 per year to reduce the violent crime rate by one percent at the mean level of violent crime (350.92 incidents per 100,000 residents).

The forward-looking model’s estimate of annual marginal willingness to pay is recovered by adjusting each household’s static estimate by the county-specific estimate of γ_k , i.e., $\hat{\alpha}_i^f = \hat{\alpha}_i^s / \hat{\gamma}_k$. This method is, by construction, equivalent to using Equation (9) directly, following the discussion in Section 3. As all values of γ_k are estimated to be strictly less than one, our forward-looking willingness-to-pay measures are larger than the static measures in absolute value. The sample mean of the distribution of these estimates is -\$14.28 with a standard error of \$0.28. That is, 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

²⁶The county-specific transition probability parameters for violent crime, i.e., $\rho_{0,k}$, $\rho_{1,k}$, and $\rho_{2,k}$, are reported in Table A.4 in the Appendix.

²⁷The 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. However, as discussed in Section 3, this trend term will not affect how estimates from our forward-looking model will differ from those from the static model.

²⁸In less than one percent of cases, $\theta_{1,k} + 2\theta_{2,k}x_{j,k,t} > 0$ implying a positive value for that household’s willingness to pay for crime exposure. We exclude these households from the calculation of utility parameters.

static approach leads to estimates that are over twenty-five percent lower in absolute terms.²⁹

Table 2: Average Annual Marginal Willingness to Pay by County

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
Static	-5.60 (0.26)	-16.54 (0.19)	-16.00 (1.47)	-4.24 (0.30)	-11.61 (0.43)
Forward-Looking	-6.40 (0.30)	-19.90 (0.22)	-31.42 (2.88)	-8.35 (0.60)	-15.55 (0.58)
Implied Bias	-12.49% (0.03)	-16.97% (0.02)	-49.09% (0.09)	-49.23% (0.05)	-25.35% (0.04)

Standard errors are given in parentheses.

Importantly, our county-level specification allows us to explore heterogeneity in the ratio of marginal-willingness-to-pay estimates. As can be seen in Table 2, 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 of those counties, there is only a small amount of mean reversion (which is the key factor 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 the violent crime rate.

Finally, as these are the same data used in Bishop and Murphy (2011), we are able to directly compare the results obtained using the simple adjustment-factor approach derived here with the fully-dynamic approach used there. As previously noted, the forward-looking model laid out in this paper is not fully dynamic, as households may not re-optimize within the T -period time horizon. The forward-looking model exogenously sets each household's probability of moving to zero for T periods (and ends the process after that), while in the fully-dynamic model, re-optimization probabilities are endogenously solved for in each period.³⁰ However, as the utility and transition-probability functions are the same in both papers, the estimates are directly comparable. Interestingly, the average forward-looking

²⁹The FBI's reported violent crime rate is a commonly-used statistic of high-profile crime in the US. However, it is likely that households respond more to homicides than to the other subcomponents of the violent crime rate. Therefore, using data describing county-level homicides for the five counties of our analysis over 1990-2008 (published by the California Department of Justice), we estimate an AR(1) process using the rate of homicide. We find mean-reversion in the homicide rate and recover an adjustment factor of $\gamma = 0.799$, with a standard error of 0.089.

³⁰While the fully-dynamic model is naturally a richer framework, solving for these endogenous re-optimization probabilities is exactly what increases the computational and data requirements.

estimate of willingness to pay found here of -\$14.28 is reasonably close to the average fully-dynamic estimate of willingness to pay found in the 2011 paper of -\$13.45.

4.2 The Willingness to Pay to Avoid Air Pollution

In this application, we revisit the well-known work of Chay and Greenstone (2005) that analyzes the housing-market impacts of the decrease in air pollution driven by the Clean Air Act (CAA) Amendments of 1970. The measure of air pollution is total suspended particulates (TSP), which was one of the primary targets of the 1970 amendments.³¹ Using a novel instrumental-variable (IV) strategy to estimate the hedonic price function for TSP, Chay and Greenstone find a large increase in housing values over the 1970s due to the CAA-driven declines in TSP.

This application is an ideal setting for us to revisit for two reasons. First, the approaches are quite complementary; the key focus of Chay and Greenstone (2005) is to credibly estimate the price function, while our focus is to develop insights regarding the mapping of these price functions into marginal-willingness-to-pay estimates. Second, we use the quasi-experimental variation associated with Chay and Greenstone to highlight the impacts of our forward-looking framework. The novel insight of Chay and Greenstone was to recognize that the CAA Amendments of 1970 generated exogenous variation in TSP concentrations that could be used to identify the price function; the 1970 Amendments subjected counties with threshold-exceeding TSP concentrations to increased regulation and deemed them to be in “non-attainment status.” This designation forms the basis of Chay and Greenstone’s IV: using an indicator for a county’s status to instrument for its decline in TSP. Interestingly, this regulatory environment generates mean-reversion in TSP concentrations, as the largest TSP declines would be seen in the counties with the highest initial concentrations.

To implement our approach, we directly use the estimate of the average implicit price of TSP found in Chay and Greenstone (2005). This estimate, which describes the causal impact of TSP on the log of housing sales price, comes from an IV, first-difference approach that includes county-level controls. Chay and Greenstone find that a one μ/m^3 (microgram per cubic meter of air) increase in TSP causes house prices to fall by 0.28%, i.e., causes the price of the average house to fall by \$243 (in 2001 dollars).

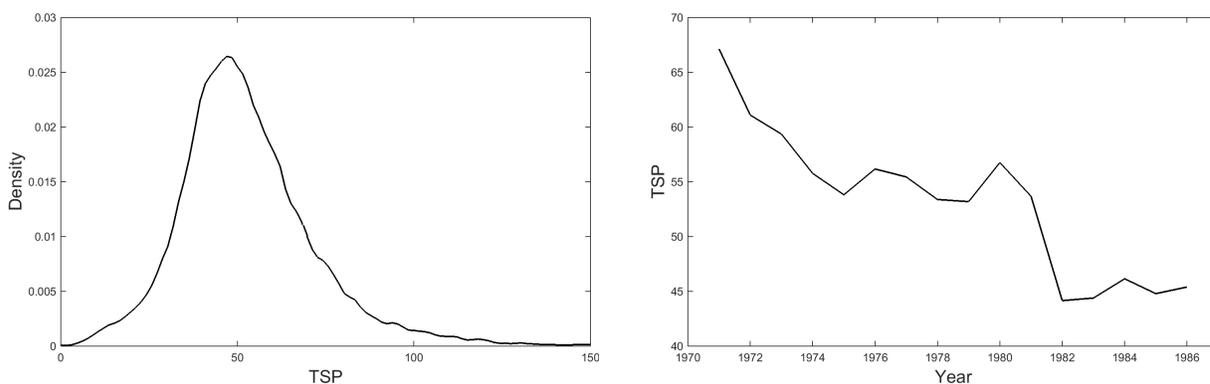
Following the discussion in Section 2, we convert this implicit price into one described by annual user costs of housing (versus sales prices) by multiplying it by 0.075, arriving at

³¹TSP includes all particulate matter that is suspended in the air, including the more-recently-regulated PM10 (suspended particulates smaller than 10 microns in size) and the currently-regulated PM2.5 (suspended particulates smaller than 2.5 microns in size).

an annualized implicit price of $-\$18.23$. Thus, our static estimate of the average marginal willingness to pay is $\hat{\alpha}^s = -\$18.23$ for a one μ/m^3 increase in TSP.³²

To estimate the transition probabilities required to form the forward-looking estimate, we use data describing county-by-year TSP concentrations from the United States EPA’s Air Quality Monitoring System over the period 1957 to 2000. These are the same underlying EPA data from which the sample used to estimate the price function in Chay and Greenstone (2005) was created. These data are organized by air-quality monitor and report annual mean TSP concentrations for each monitor in the sample, which we aggregate to county level.³³

Figure 5: Variation in County-Level TSP Concentration (μ/m^3)



While we have access to forty-four years of data, we create a sixteen-year sample of 1971 to 1986. This period reflects the post-CAA years over which TSP was the target of EPA regulation. The 1970 Amendments began implementation in 1971. Beginning in 1987, the EPA stopped using TSP as its preferred measure of particulates and many monitors ceased to record TSP concentrations.³⁴ There is a significant downward trend in TSP concentrations over this time period. We show this time-series variation, along with the distribution of TSP, in Figure 5.

We specify the transition of TSP as following an AR(1) process:

$$x_{k,t} = \rho_0 + \rho_1 x_{k,t-1} + \rho_2 t + \varepsilon_{k,t} \tag{19}$$

³²By multiplying the implicit price by 0.05, Chay and Greenstone (2005) report a somewhat analogous number of $0.05 \times \$243 \approx \12 . Previewing the importance of calculating a forward-looking estimate, they state: “Of course, the valuation calculations for a one-unit reduction in TSPs will vary depending on the discount rate and individuals expectations on the future path of TSPs.”

³³To arrive at this annual mean TSP concentration, the EPA takes a geometric mean of each monitor’s daily TSP concentrations in a year. To aggregate to the county level, we follow Chay and Greenstone and take a weighted (arithmetic) mean over all monitors’ annual means within each county. The number of each monitor’s readings serve as weights.

³⁴See Kahn (1997) for a detailed discussion of the timeline.

where k denotes county and t denotes year.³⁵ As in our previous application, this simple transition process allows us to express \bar{x} as $\bar{x}_{k,t} = \phi_t + \gamma x_{k,t}$, where γ (the term that determines the bias) is given by the weighted average

$$\gamma = \frac{\sum_{s=1}^T \beta^{s-1} \rho_1^{s-1}}{\sum_{s=1}^T \beta^{s-1}} \quad (20)$$

and, as before, β is set to 0.95 and T is set to seven years.

Our estimate of γ is 0.519 (with a standard error of 0.005).³⁶ This indicates a significant level of mean reversion, as would be expected given the CAA-motivation. This also indicates that a large adjustment would need to be made to the static estimate of willingness to pay. In particular, we construct the forward-looking estimate of average marginal willingness to pay for a one μ/m^3 increase in TSP by dividing the static estimate of $\hat{\alpha}^s = -\$18.23$ by $\hat{\gamma} = 0.519$. Thus, our forward-looking estimate of average marginal willingness to pay is $\hat{\alpha}^f = -\$35.14$, or almost double the static-model estimate.

To allow for the possibility that the CAA itself affected the expectation-formation process, we also estimate our model using the pre-CAA years of 1957 to 1970. Using this pre-CAA sample, the estimate of γ is 0.640 (with a standard error of 0.016). This estimate, which indicates a lower level of mean reversion prior to 1971, is to be expected; as carefully documented in Chay and Greenstone (2005), the CAA Amendments provided much stronger incentives for dirty counties to reduce their TSP concentrations, thus contributing to the higher levels of mean-reversion post-CAA. This estimate of γ implies a forward-looking estimate of marginal willingness to pay of $\hat{\alpha}^f = -\$28.48$.

To conclude, Chay and Greenstone (2005) showed that by using credible, exogenous variation in TSP (driven by the CAA Amendments of 1970), the housing-market impacts of air pollution were larger than previously thought, as existing studies recovered small and/or statistically insignificant impacts of TSP in the housing market.³⁷ We build on that conclusion by further showing that, due to the mean-reverting nature of the variation in TSP, the implications for household preferences would be even larger again.

³⁵We follow Chay and Greenstone (2005) and treat the U.S. as a national market.

³⁶This estimate is robust to alternative specifications of the formation of expectations. Specifying an AR(2) process yields an estimate of γ of 0.541, with a standard error of 0.005. Using data describing all available TSP readings from 1971 until 2000 yields an estimate of γ of 0.543, with a standard error of 0.004.

³⁷Chay and Greenstone themselves present three “conventional” estimates of the housing market impacts of TSP reductions: estimates from a 1970 cross-sectional regression (small but significant results), estimates from a 1980 cross-sectional regression (counter-intuitively signed but significant results), and estimates from a 1973-1980 first-difference regression with no IV (small and insignificant results).

5 Conclusion

Researchers in a wide variety of applied fields have relied on Rosen’s intuitive 1974 model to recover households’ 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 marginal willingness-to-pay estimate from the static model into one from the forward-looking model.

We highlight these concepts with two empirical applications. In the first, we recover households’ willingness to pay to avoid violent crime in each of five counties in the Bay Area of California. We find considerable heterogeneity across counties in the rates of mean-reversion of violent crime and, therefore, considerable heterogeneity in the bias associated with specifying the traditional, static model. In the second application, we recover households’ willingness to pay to avoid air pollution. With significant mean-reversion driven by the Clean Air Act Amendments of 1970, we find a large difference between the estimates of marginal willingness to pay from the static and forward-looking models. The results of both of these applications 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 assess the potential benefits prior to adopting either a static or 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., J. C. FRUEHWIRTH, K. I. KIM, AND C. TIMMINS (2012): “A rational expectations approach to hedonic price regressions with time-varying unobserved product attributes: The price of pollution,” *The American Economic Review*, 102(5), 1898–1926.
- 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 (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. 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. (2016): “Neighborhood Sorting and the Value of Public School Quality,” mimeo, University of Rochester.
- CELLINI, S. R., F. FERREIRA, AND J. ROTHSTEIN (2010): “The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design,” *The Quarterly Journal of Economics*, 125(1), 215.
- CHALFIN, A., AND J. MCCRARY (2015): “Are U.S. Cities Underpoliced?: Theory and Evidence,” mimeo.
- 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 Functional 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 (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.
- 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.

KAHN, M. E. (1997): “Particulate Pollution Trends in the United States,” *Regional Science and Urban Economics*, 27(1), 87–107.

KANE, T. J., AND D. O. STAIGER (2008): “Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation,” NBER Working Paper No. 14607.

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.

MENDELSON, 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 General Framework

We make the assumption that utility is increasing and concave in the amenity x , but make no other assumptions on its functional form.

In the static model, households choose x_i to maximize $U(x_i)$, where $U(x_i)$ is given by:

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

This yields the first-order condition:

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

which we can then re-write as:

$$u'(x) = r'(x) \quad (\text{A.3})$$

In the forward-looking model, households choose $x_{i,t}$ to maximize $v(x_{i,t})$, where $v(x_{i,t})$ is given by:

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

This yields the first-order condition:

$$v'(x_{i,t}) = \sum_{s=1}^T \beta^{s-1} \frac{\partial E[u(x_{i,t+s-1}) \mid x_{i,t}]}{\partial x_{i,t}} - \sum_{s=1}^T \beta^{s-1} r'(x_{i,t}) = 0 \quad (\text{A.5})$$

which we can then rewrite as:

$$\frac{1}{B} \sum_{s=1}^T \beta^{s-1} \frac{\partial E[u(x_{i,t+s-1}) \mid x_{i,t}]}{\partial x_{i,t}} = r'(x_{i,t}) \quad (\text{A.6})$$

where $B = \sum_{s=1}^T \beta^{s-1}$.

As may be seen in the comparison of the left-hand sides of Equations (A.3) and (A.6), the bias is determined by whether $\frac{1}{B} \sum_{s=1}^T \beta^{s-1} \frac{\partial E[u(x_{i,t+s-1}) \mid x_{i,t}]}{\partial x_{i,t}} = \frac{\partial u(x_{i,t})}{\partial x_{i,t}}$.

First, we note that the intuition developed in Section 3.2 for mean-reverting and mean-

diverging cases still holds.³⁸ Second, while most of the intuition developed in Section 3.1 for increasing and decreasing amenities continues to hold, the specific conclusion that there is no difference between the static and forward-looking estimates, does not; the price and quantity effects still work in opposite directions, but do not *exactly* offset one another when the utility function is non-linear.

As the intuition for the mean reverting/diverging case still holds, we provide a sketch of the direction of the bias under the alternative simple trends of the amenity (in the absence of mean-reversion or diversion):

Case 1: There is no trend in the amenity – $E[x_{i,t+s-1}|x_{i,t}] = x_{i,t}$. As utility is concave, Jensen’s Inequality means that $E[u(x_{i,t+s-1})|x_{i,t}] < u(E[x_{i,t+s-1}|x_{i,t}])$. 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 $x_{i,t}$ than is $u(E[x_{i,t+s-1}|x_{i,t}])$. As $E[x_{i,t+s-1}|x_{i,t}] = x_{i,t}$, then $E[u(x_{i,t+s-1})|x_{i,t}] < u(x_{i,t})$ and $\frac{\partial E[u(x_{i,t+s-1})|x_{i,t}]}{\partial x_{i,t}} > \frac{\partial u(x_{i,t})}{\partial x_{i,t}}$. In other words, the static model overstates marginal willingness to pay.

Case 2: There is a decreasing trend in the amenity – $E[x_{i,t+s-1}|x_{i,t}] < x_{i,t}$. In this case, in addition to the effect described in Case 1, the concavity of utility will lead $E[u(x_{i,t+s-1})|x_{i,t}]$ to be steeper than $u(x_{i,t})$. In other words, the static model unambiguously overstates marginal willingness to pay.

Case 3: There is an increasing trend in the amenity – $E[x_{i,t+s-1}|x_{i,t}] > x_{i,t}$. In this case, the concavity of utility will lead $E[u(x_{i,t+s-1})|x_{i,t}]$ to be less steep than $u(x_{i,t})$. On its own, this would lead the static model of understate marginal willingness to pay. However, the effects described in Case 1 also hold. Depending on which effect dominates, the static model will either under- or over-estimate marginal willingness to pay. In general, the faster the rate of the increase in the amenity, the more likely it will be that the static model understates marginal willingness to pay.

While this intuition for the direction of the bias can be developed for the non-linear case, the simple adjustment factor cannot be recovered for all forms of non-linear utility. However, an adjustment factor can easily be derived for common forms of non-linear utility.³⁹ As an example, we outline the commonly-used log-utility case.

³⁸It is worth noting that as the household observes current values of the amenity, the variance that affects the household’s behavior is the conditional variance, $\text{var}(x_{i,t+1}|x_{i,t})$, rather than the unconditional variance, $\text{var}(x_{i,t+1})$. Therefore, even though the unconditional variance, $\text{var}(x_{i,t+1})$, is falling over time in the mean reversion case, this does not imply anything about the conditional variance. In our empirical application, the conditional variance, $\text{var}(x_{i,t+1}|x_{i,t})$, is given by $\text{var}(\varepsilon_{j,k,t})$, where $\varepsilon_{j,k,t}$ is defined in Equation (A.12).

³⁹An adjustment factor can be derived for the set of models described by a single utility parameter, which corresponds to the set of static models that can be estimated using the familiar Bajari and Benkard (2005) framework.

A.2 Example with Log Utility

We consider the following model

$$U(x_i) = \alpha_i \log(x_{i,t}) + I_i - r(x_{i,t}) \quad (\text{A.7})$$

which yields a per-period marginal willingness to pay of $\alpha_i/x_{i,t}$.

Analogous to the linear case, we define a measure of expected average $\log(x)$ consumption over the horizon T with the following weighted average:

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

Equation (A.6) can then be written as:

$$\alpha_i \frac{\partial \bar{x}_{i,t}}{\partial x_{i,t}} = r'(x_{i,t}) \quad (\text{A.8})$$

which (following some simple algebra) can then be written as:

$$\alpha_i \frac{\partial \bar{x}_{i,t}}{\partial \log(x_{i,t})} \frac{\partial \log(x_{i,t})}{\partial x_{i,t}} = r'(x_{i,t}) \quad (\text{A.9})$$

$$\alpha_i \frac{\partial \bar{x}_{i,t}}{\partial \log(x_{i,t})} \frac{1}{x_{i,t}} = r'(x_{i,t}) \quad (\text{A.10})$$

An estimate of the per-period marginal willingness to pay can be constructed as:

$$\frac{\hat{\alpha}_i^f}{x_{i,t}} = \frac{1}{\hat{\gamma}} \hat{r}'(x_{i,t}) \quad (\text{A.11})$$

where $\gamma = \frac{\partial \bar{x}_{i,t}}{\partial \log(x_{i,t})}$.

While, in theory, γ could be recovered from a specification where x follows an autoregressive process, the problem is much simpler if $\log(x)$ follows an autoregressive process such as:⁴⁰

$$\log(x_{j,t}) = \varrho_0 + \varrho_1 \log(x_{j,t-1}) + \varrho_2 t + \varepsilon_{j,t} \quad (\text{A.12})$$

⁴⁰By adopting a AR(1) model for the log-utility case, we are changing the functional form of expectations as well as the functional form of utility. However, the simple nature of how the adjustment factor is constructed is maintained.

and γ is given by:

$$\gamma = \frac{\sum_{s=1}^T \beta^{s-1} \varrho_1^{s-1}}{\sum_{s=1}^T \beta^{s-1}} \quad (\text{A.13})$$

In practice, therefore, the only difference between recovering the adjustment factor in the log-utility versus linear-utility case is whether one specifies the transition process in logs or levels. By definition, the static-model estimate of marginal willingness to pay (at the point of consumption) is the same in the log-utility and linear-utility cases.

Table A.1: Average Marginal Willingness to Pay by County

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
linear utility (baseline)	-6.40 (0.30)	-19.90 (0.22)	-31.42 (2.88)	-8.35 (0.60)	-15.55 (0.58)
log utility	-6.53 (0.31)	-20.79 (0.23)	-31.63 (2.90)	-9.83 (0.70)	-15.58 (0.58)

Standard errors are given in parentheses.

We empirically calculate the log-utility adjustment factor for our violent crime application and show in Table A.1 that the estimates of marginal willingness to pay are very similar in the linear-utility and log-utility cases.⁴¹ The minor differences are consistent with the logic developed in Section A.1: with falling crime (i.e., increasing safety), the log-utility estimates are bigger (in absolute value) and the biggest difference occurs in San Mateo, where crime is falling the fastest.

⁴¹As violent crime is a “bad” and not a “good”, we specify $x = \text{constant} - \text{violent crime}$. The constant is chosen to be 1.1 times the maximum observed level of violent crime to ensure that x , which is a measure of safety, is positive. See Bishop and Timmins (2017) for a similar setup with air pollution.

Appendix B: Additional Data Details

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



Many of the sample cuts are made in order to deal with the fact that we only see the housing characteristics at the time of the property’s 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 municipality-level crime rates using each property’s geographic coordinates, we drop properties where latitude and longitude are missing. The centroids of the 75 municipalities are shown in Figure A.1.

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 violent 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. Summary statistics are shown in Tables A.2 and A.3.

Table A.2: Property Transactions Summary Statistics

<i>N = 541,415</i>					
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 sq. footage	1,687	1545	662	160	9,130
lot Sq. 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.3: Household Summary Statistics

<i>N = 372,334</i>					
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: Additional Results

Table A.4: Transition Probability Estimates

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
constant	30.4772 (0.1345)	17.1805 (0.1128)	110.6628 (0.3849)	94.5984 (0.3438)	49.5586 (0.0987)
lagged violent crime rate	0.9518 (0.0001)	0.9329 (0.0001)	0.7387 (0.0008)	0.7375 (0.0004)	0.8926 (0.0002)
t	-1.2274 (0.0083)	0.0473 (0.0080)	-3.8823 (0.0188)	-0.5867 (0.0226)	-2.3091 (0.0048)
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

These are the OLS estimates of Equation (17). Standard errors are given in parentheses.

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

	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
1990	67.27 (0.27)	43.60 (0.22)	179.65 (0.41)	173.12 (0.44)	96.70 (0.17)
t	-3.18 (0.02)	0.12 (0.02)	-7.29 (0.03)	-1.10 (0.04)	-5.45 (0.01)

These are the remaining estimates describing the following relationship:

$\bar{x}_{j,k,t} = \phi_{k,t} + \gamma_k x_{j,k,t}$. Standard errors are given in parentheses.

Table A.6: Hedonic Price Function Estimates

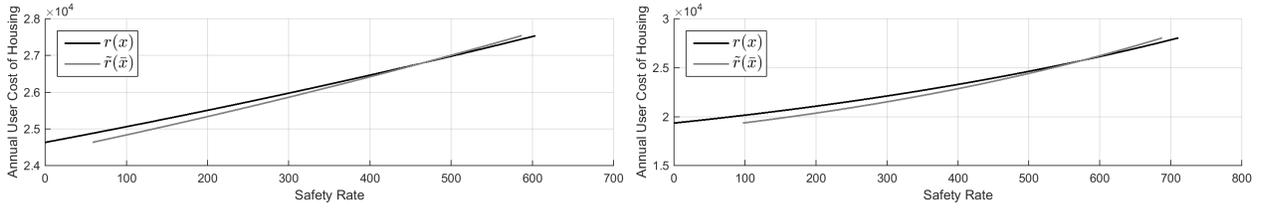
	Alameda	Contra Costa	Marin	San Mateo	Santa Clara
violent crime rate	-0.20903 (0.01361)	-0.70398 (0.00803)	-0.96640 (0.08403)	-0.12529 (0.01014)	-0.46180 (0.01977)
violent crime rate squared	0.02600 (0.00660)	0.18787 (0.00274)	1.26898 (0.11632)	0.02091 (0.00357)	0.22928 (0.01556)
house sq. footage	0.29287 (0.00140)	0.29286 (0.00154)	0.37959 (0.00259)	0.22190 (0.00230)	0.30755 (0.00153)
lot sq. footage	0.00862 (0.00012)	0.00826 (0.00008)	0.00418 (0.00018)	0.00865 (0.00014)	0.00654 (0.00007)
house age	-0.00151 (0.00004)	-0.00133 (0.00005)	0.00090 (0.00009)	0.00076 (0.00005)	0.00054 (0.00003)
number of rooms	0.01839 (0.00039)	0.03168 (0.00050)	0.00143 (0.00055)	0.05140 (0.00079)	0.03693 (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

These are the OLS estimates from the estimation of Equation (16). Standard errors are given in parentheses.

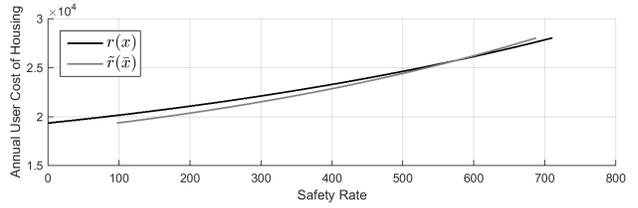
The variables of violent crime rate, violent crime rate squared, house sq. footage, and lot sq. footage are measured in 1000s of units.

For the figures presented in Figure A.2, we plot the estimates of the housing price regressions shown in Table A.6 (the precision of the underlying estimates may be seen in Table A.6). For each county, we plot the housing price function, $r(x)$, on the axes above and the implicit price function, $r'(x)$, on the axes below. To keep these figures consistent with our earlier theoretical framework (i.e., located in the first quadrant and describing a “good”), we plot the annual user cost of housing against safety, where the safety rate is defined as a county-specific constant minus the violent crime rate, such that values of safety are positive. As the functions in Figure A.2 are plotted from the 5th to 95th percentiles of county-specific safety rates (holding each of the control variables in $H_{j,k,t}$ at their means), the domains display both the overall change and mean reversion in safety for each county.

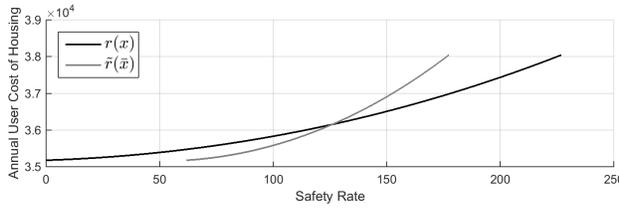
Figure A.2: Estimated Price Functions and Hedonic Gradients for Each County



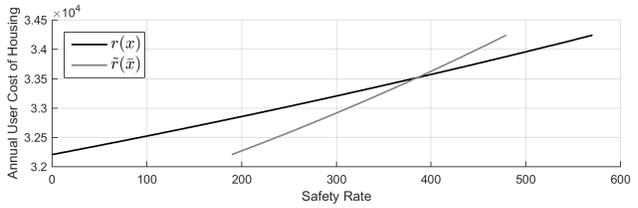
(a) Alameda County



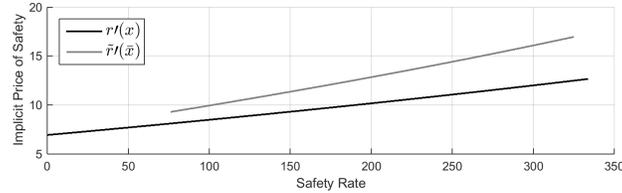
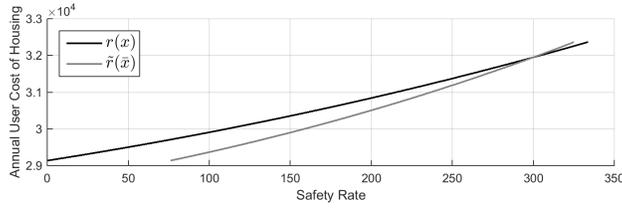
(b) Contra Costa County



(c) Marin County



(d) San Mateo County



(e) Santa Clara County

Appendix D: Sensitivity Analysis

In Table A.7, we provide the results from a variety of robustness checks to our forward-looking estimates of average willingness to pay to avoid one additional violent crime per 100,000 local residents. Row I presents the baseline results discussed in the text, where violent crime follows an AR(1) process, $T = 7$, and $\beta = 0.95$.

Table A.7: Average Marginal Willingness to Pay by County

		Alameda	Contra Costa	Marin	San Mateo	Santa Clara
I.	baseline	-6.40 (0.30)	-19.90 (0.22)	-31.42 (2.88)	-8.35 (0.60)	-15.55 (0.58)
II.	AR(2)	-6.61 (0.31)	-19.74 (0.22)	-27.55 (2.52)	-7.08 (0.51)	-15.38 (0.57)
III.	adaptive expectations	-6.32 (0.30)	-20.00 (0.23)	-32.82 (3.01)	-8.87 (0.63)	-14.93 (0.56)
IV.	air pollution control	-6.25 (0.29)	-19.78 (0.22)	-31.60 (2.89)	-8.45 (0.60)	-15.39 (0.58)
V.	$T = 5$	-6.13 (0.29)	-18.78 (0.21)	-26.01 (2.38)	-6.91 (0.49)	-14.22 (0.53)
VI.	$T = 9$	-6.66 (0.31)	-21.00 (0.24)	-36.80 (3.37)	-9.78 (0.70)	-16.88 (0.63)
VII.	$\beta = 0.90$	-6.33 (0.30)	-19.61 (0.22)	-29.68 (2.72)	-7.88 (0.56)	-15.19 (0.57)
VIII.	$\beta = 1$	-6.46 (0.31)	-20.18 (0.23)	-33.25 (3.05)	-8.84 (0.63)	-15.91 (0.59)

Standard errors are given in parentheses.

Rows II and III present estimates under alternative specifications for the formation of expectations. Under the assumption of rational expectations, households know the underlying transition process of violent crime. In Row II, we maintain the baseline assumption of rational expectations, but allow violent crime to follow an AR(2) process. As can be seen, results are similar. In Row III, we allow for a departure from rational expectations and assume that households do not know the common, underlying transition process, but have adaptive expectations that update each year. To implement this assumption, we estimate a separate AR(1) for each year that employs only previous years of data. Again, results are similar.

In Row IV, we present estimates from an alternative specification for the transition of violent crime that includes air pollution (in the form of the contemporaneously-federally-regulated PM10) as an additional control variable. Results are similar.

Rows V and VI present estimates under alternative time horizons. In Row IV, we shorten each household's horizon to five years. In light of the fact that $T = 1$ would collapse to the static case, the results in Row IV are similar to those in Row I. In Row V, we extend each household's horizon to nine years. Again, results are quite similar.

In Rows VII and VIII, we present estimates under alternative discount rates, $\beta = 0.90$ and $\beta = 1$, respectively. In both cases, results are similar to those in Row I.