Connect-The-Dots: Identification of Heterogeneous Marginal Willingness to Pay Functions under Time-Varying Preferences

Maria Juul Hansen*

Abstract

Models based on residential location choice have become commonplace in the non-market valuation literature. Rosen (1974) provides a utility-theoretic basis for hedonic models to be used to measure the welfare consequences of changes in local public goods and amenities. However, his proposed two-stage estimation procedure embodies a number of difficult econometric problems that have become the focus of research for decades. My paper builds upon the "inversion" approach suggested by Bajari and Benkard (2005) and the buyer-panel extension of that work proposed by Bishop and Timmins (2018). The latter paper shows how data on repeat purchases can be used to flexibly recover preferences with rich individual heterogeneity, but the method is unable to deal well with time-varying individual attributes that might prompt residential location changes. I expand that approach to deal with any number of time-varying individual attributes including income, family structure and other drivers of housing choice. I apply that method to detailed longitudinal data from the Danish census, and use the estimates to value non-marginal changes in violent crime rates. I demonstrate a significant and policy-relevant bias from failing to properly account for the endogeneity problems in Rosen (1974).

Keywords: Non-Market Valuation, Marginal Willingness to Pay, Hedonic Analysis, Violent Crime

JEL classification: Q50, Q51, R21, R23

*University of Copenhagen, maria.juul.hansen@econ.ku.dk

1 Introduction

The valuation of neighborhood amenities and local public goods is important for the allocation of public funds and the measurement of the benefits of regulation and other policies. The value of these goods cannot generally be measured from market prices, but because many of them are, as their name suggests, local (i.e., consumption varies with geography), their values can be recovered from residential choices. Rosen (1974) provides the theory that connects those decisions to utility-theoretic measures of welfare, setting the stage for hedonic theory to be used in a variety of policy contexts.

While it is the basis for an entire literature, Rosen's procedure for recovering preferences from housing decisions is problematic. By the 1980s it was realized that there were important endogeneity problems inherent in his approach (Epple, 1987; Bartik, 1987), and the literature began to suggest alternatives. This paper builds upon an "inversion" approach suggested by Bajari and Benkard (2005), in which preferences are not estimated in a traditional sense, but rather they are recovered at an individual level from the conditions imposed by optimizing behavior. This avoids the need for an unobserved (to the econometrician) preference shock, which is the source of the econometric problems mentioned above.

Specifically, my approach extends the idea in Bajari and Benkard (2005) to include the information available in a buyer-panel – individuals who are observed to buy more than one housing unit over the span of many years. This repeat-buyer information allows one to more richly describe individual preferences than was feasible using the cross-sectional individual information about consumers described by Bajari and Benkard (2005). Bishop and Timmins (2018) make this point and demonstrate the power of a buyer-panel in this context. However, their paper also reveals the weakness inherent in buyer-panel data. Specifically, it takes many years to observe individuals buying multiple houses, and during that time horizon one would expect many of their circumstances to change. Indeed, it is often a change in household circumstances that prompts a move. Changing circumstances implies a time-varying form of individual heterogeneity that may be hard to observe. Panel data sets in the U.S. generally either (i) are rich in their description of individuals but available at a level of spatial disaggregation that is insufficient for modeling exposure to local public goods and amenities (e.g., the PSID or NLSY), or (ii) do not contain information about salient household attributes (e.g., data constructed from housing transaction information linked by buyers' names)[1]. I overcome this data constraint by employing restricted-access census data from Denmark. These data provide detailed information about numerous static and time-varying household attributes, in addition to specific information about the location and value of purchased homes. These data allow the researcher to control for family structure (e.g., marriage, divorce, death, or the birth of a new child), changes in income or wealth, education, and other potentially important characteristics. Bishop and Timmins (2018) ignored nearly all of these time-varying drivers of demand for amenities because the data needed to address them were unavailable. Moreover, with the theory in that paper, it is difficult

[1] Bishop and Timmins (2018) use data on housing transactions from a real estate data services provider, linked by buyer name and sale/purchase dates. These data can be linked to information about race and income collected under the Home Mortgage Disclosure Act, but that is the extent of household heterogeneity that can be included.

to address time-varying observables even with the requisite data. This is because adding more time-varying attributes requires adding more repeat purchases to achieve identification, meaning a longer panel is required, and this generally means that more changing household attributes are likely to be relevant. In this paper, I break that cycle with an alternative technique that uses the large size of the Danish census to adjust characteristics such that I can argue I observe the individual on the same demand curve twice when she buys a new property. With two points on the same demand curve I am able to identify her marginal willingness to pay (MWTP) for local amenities in a flexible way that avoids the well-known problems with Rosen's procedure i.e., by simply connecting the dots.

I apply my method to an array of local public goods and amenities that might determine residential location choice. I focus my attention on violent crime, showing how the value of an infra-marginal change differs dramatically depending upon whether one recovers an unbiased measure of preferences in the form of a MWTP function or not. My results suggest that ignoring the endogeneity problems of Rosen's procedure means understating the welfare costs of a large (i.e., 30 percent) increase in violent crime by up to 70% and 24.4% on average. I also show there is a high degree of heterogeneity in the willingness to pay for avoiding increases in violent crime across the violent crime distribution, and that the bias is even more severe for individuals living in high-crime areas. Targeting resources based on biased estimates of MWTP therefore means the spatial allocation of these resources will be highly distorted; low-crime areas will receive an amount of resources that reflects their WTP to a higher degree (though still not exactly) than the high-crime areas will. From an economic viewpoint, policy makers should target amenity-improving resources to locations where there is a high WTP for such improvements and the approach in this paper therefore has huge potential.

The paper proceeds as follows. Section 2 describes in more detail the source of the difficulty in recovering an unbiased estimate of the MWTP function. Section 3 lays out the theory of my proposed estimation strategy, while Section 4 discusses my data, which are particularly well-suited to modeling panel variation in house purchase decisions. Section 5 reports results with a particular emphasis on the differences in estimated value of willingness to pay for a large change in violent crime using the alternative estimation strategies. Section 6 concludes.

2 The difficulty with estimating MWTP

In his 1974 article, Rosen proposed a two-step estimator for the MWTP function. His insight was the following: The slope of an individual's indifference curves in (q, P) space, where q represents some amenity and P is the price of the house associated with that amenity, reflects the willingness to give up additional units of other consumption (in the form of paying more for a house) in exchange for more q. Conveniently, individuals will sort into the housing unit that maximizes their utility, and that point on the slope of the hedonic price function will reveal the slope of their indifference curve. Figure 1 illustrates how agents A and B optimize where their indifference curves U^A and U^B are tangent to the hedonic price function P.

Rosen proposed a two-step approach to recovering preferences from the hedonic price func-

Figure 1: Picking q to optimize utility

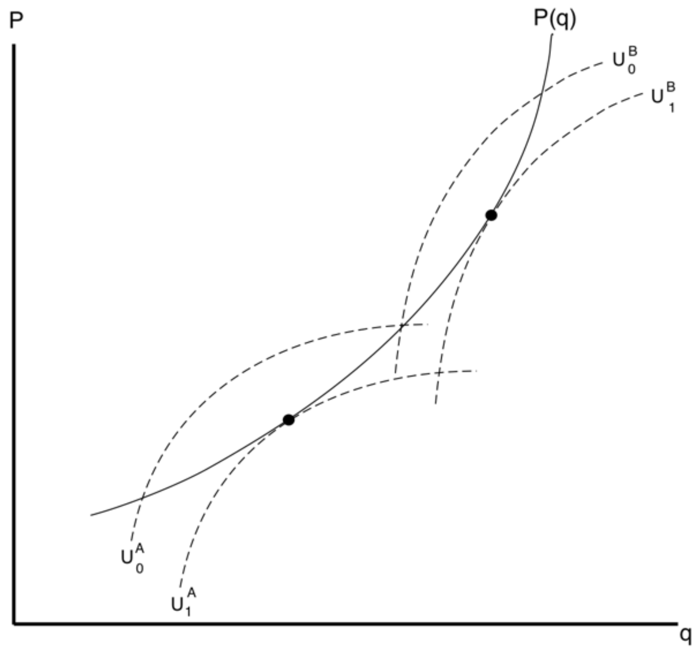

tion. In the first step, the hedonic relationship between price (P), the amenity (q) and other housing characteristics (x) is estimated. For each observed house purchase i, the implicit price of q ($\partial P_i/\partial q_i$) is calculated, and is then used as the dependent variable in an estimation of the demand for q:

$$P_i^q = \frac{\partial P_i}{\partial q_i} = \gamma_0 + \gamma_1 q_i + \gamma_2 w_i + \epsilon_i \tag{1}$$

where w_i is a vector of individual attributes[2]. Rosen suggested estimating this equation by OLS and using the result to measure the value of a large change in the amenity q by integrating to yield measures of consumer surplus[3].

The literature eventually pointed out two potential problems with this approach. The first is that if P_i^q is just a function of q_i, there is no additional information introduced by the hedonic gradient. This exercise then amounts to regressing a linear function of q_i on q_i, and Brown and Rosen (1982) show that this will simply reproduce the hedonic gradient. Brown and Rosen (1982) and Mendelsohn (1985) show how this problem can be addressed by imposing functional form restrictions on the hedonic price and MWTP functions, or by exploiting data from multiple markets. More recently, Ekeland et al. (2004) show that the type of linearity that causes the problems noted by Brown and Rosen (1982) and Mendelsohn (1985) is a special case and need not be a concern.

A second form of endogeneity was noted by Epple (1987) and Bartik (1987). They pointed out that, when an individual sorts along the hedonic price function, she both chooses the level

[2]Note that it is relatively easy to observe detailed information about households' characteristics if one is able to use cross-sectional data for estimation, as is required for Rosen's method.

[3]Willig (1976) discusses the role of income effects and the difference between compensated and uncompensated demand for welfare analysis.

of the amenity (q_i) and the implicit price that she pays for it (P_i^q). If the hedonic price function is non-linear, in the notation of the previous equation, individuals with large values of ϵ_i (i.e., strong preferences for q) will choose a high value of q (q_B in Figure 2) and, if the hedonic price function is convex, a high value for the implicit price of q (P_B^q in Figure 2). ϵ_i is thus correlated with both q_i and P_i^q. Epple (1987) notes that the traditional approach to using instruments in a system of supply and demand equations will not work here because buyers and sellers are systematically matched with one another by the sorting process. The literature struggled with this problem for more than a decade.

Figure 2: Picking q to optimize utility with a non-linear hedonic price function

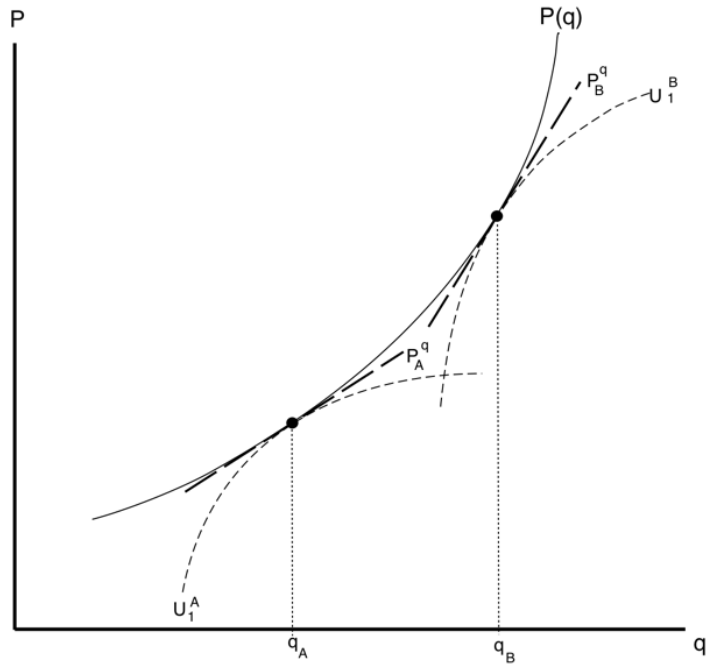

Bajari and Benkard (2005) provide a solution to this problem by replacing the estimation approach suggested by Rosen (1974) with an inversion approach. In particular, they propose writing down a utility function with heterogeneity embodied in utility function parameters at the individual level. For the sake of simplicity, I use a linear utility function for exposition:

$$U(q, x; \kappa) = \kappa_{1,i} q + \kappa_{2,i} x + c \tag{2}$$

which is maximized subject to a budget constraint:

$$c + P(q, x) = I, \tag{3}$$

where q is an amenity, x is other house characteristics and c is numeraire consumption and I is total income. The hedonic price function $P(q, x)$ represents the equilibrium of interactions between housing buyers and sellers and is assumed to be continuous and dense in attribute space. Solving for indirect utility (V) and taking first-order conditions with respect to q and x, one gets

$$\frac{\partial V_i}{\partial q} : \kappa_{1,i} - \frac{\partial P}{\partial q} = 0 \tag{4}$$

$$\frac{\partial V_i}{\partial x} : \kappa_{2,i} - \frac{\partial P}{\partial x} = 0. \tag{5}$$

These two equations can then be easily solved for values of $(\kappa_{1,i}, \kappa_{2,i})$ for each individual. The attractiveness of this approach comes in that individual heterogeneity is captured directly by utility function parameters, rather than in an econometric error term. It therefore avoids the endogeneity problems that accompany the latter. The downside to this approach is that it puts strong constraints on the shape of the MWTP function. In the simple example described above, those MWTP functions are necessarily horizontal lines (i.e., elasticity of zero). By assuming a Cobb-Douglas utility function, one assumes a MWTP function with elasticity of -1. If the goal is to let the data reveal the elasticity of demand for q, this is a severe limitation.

Bishop and Timmins (2018) show how the set of MWTP parameters one can recover from this inversion procedure can be expanded to include both intercept and slope parameters if one is able to observe two purchases by each individual. Importantly, that approach requires that both purchases lie on the same demand curve (i.e., preferences do not shift between house purchases). That paper also shows that identifying the coefficients on time-varying determinants of preferences requires additional repeat sales data (e.g., identifying a linear MWTP with one time-varying preference shifter would require data on three house purchases during which time no other individual attributes could vary). This creates a vicious cycle, whereby more housing transactions are required to identify the effects of time-varying preference shifters, but by including more housing transactions one increases the time dimension of the panel and number of other time-varying attributes that might change. In this paper, I demonstrate a way to break out of this cycle by using a function relating consumption of q to individual attributes to adjust that consumption to a counterfactual that holds time-varying individual attributes fixed. This approach relies on having rich data describing individual characteristics, which I get from the Danish census.

3 Using panel data to estimate MWTP

Below I outline an approach that solves the problem of identifying individuals' preferences for an amenity when the preferences are changing over time. To set the stage, I aim at estimating heterogeneous linear[4] marginal willingness to pay curves for an amenity q with parameters $\mu_i = (\mu_{i0}, \mu_{i1})$.

$$MWTP_{q_{it}} = \mu_{i0} + \mu_{i1} q_{it} \tag{6}$$

Linear functions are identified whenever two points on the line are observed. Bishop and Timmins (2018) suggest exploiting that some individuals buy a home twice and therefore reveal their demand for q in two different markets (time periods) under, potentially, two different price schedules. However, these two points would only identify μ_i as long as the individual's preferences are unchanged between the first and second purchase. Since individuals' preferences may

[4]To identify non-linear marginal willingness to pay functions, one would need a panel of buyers who are observed buying more than twice during the sample period.

change throughout their lives and it may be these changing preferences that make them decide on buying a new home, one is not guaranteed that their chosen quantities of q are actually observed along the same demand curve.

To circumvent this problem, I suggest using information on the multiple-buyer individuals *in addition to* knowledge about their potentially time-varying background characteristics to predict the level for q that the individual would have chosen had her attributes not changed. That allows the researcher to estimate the intercept and slope using the consumption of q in each period and the associated implicit prices. Below I go over each of the steps in the process of identifying the MWTP function parameters using this methodology.

3.1 Step 1: Hedonic gradient

The first step in my model is standard and is about recovering an estimate of the hedonic price function, P_{it}. The model does not rely on any particular assumptions about the shape of the hedonic price, but is generally considered a function of the amenity q_{it}, housing and neighborhood attributes, x_{it}, and some set of parameters β_t (that potentially depends on time) through the unknown function $g(.)$:

$$P_{it} = g(q_{it}, x_{it}; \beta_t). \tag{7}$$

Knowing P_{it}, the implicit price of q_{it}, P^q_{it}, is given by the derivative of the price function:

$$\frac{dP_{it}}{dq_{it}} \equiv P^q_{it} \tag{8}$$

For any given level of q_{it} the researcher is then in the position to get an estimate of the implicit price.

3.2 Step 2: Segmentation equations

To predict the quantity of violent crime at the time of the second purchase were the individual's attributes to not have changed, I propose estimating the relationship between individual demographic characteristics and the observed demand for the amenity The relationship between the quantity of q consumed and the attributes of the individuals doing that consumption describes how the market is segmented based on individual attributes. This function arises from sorting in hedonic equilibrium as illustrated by Mendelsohn (1985). It can be generally defined as

$$q_{it} = f(w_{it}; \delta_t), \tag{9}$$

where $f(.)$ is some function of w_{it}, a vector of individual demographics that are expected to affect the preferences for q, and δ_t, a vector of parameters that may depend on time.

With the segmentation equation estimated, one can then adjust the individual's second purchase characteristics back to the values she had at the time of the first purchase and label these adjusted characteristics $\tilde{w}_{i2}$. By using the estimates (denoted by hat) from the segmentation equation one can predict the adjusted demand at the time of the second purchase denoted $\tilde{q}_{i2}$:

$$\tilde{q}_{i2} = f(\tilde{w}_{i2}; \hat{\delta}_2). \tag{10}$$

That is, $\tilde{q}_{i2}$ is the estimate of the counterfactual optimal choice in period two if the individual were to still have her attributes from period one. Demand at the first purchase, q_{i1}, is directly observed in the data together with the price P_{i1}, but these can also be predicted using the estimates of the segmentation equation and gradients.

That is exactly what I suggest doing for the counterfactual period-two demand; compute the counterfactual implicit price that i would have had to pay for her counterfactual quantity $\tilde{q}_{i2}$ in period two by evaluating the estimated hedonic gradient for that period:

$$\tilde{P}^{\tilde{q}}_{i2} = g(\tilde{q}_{i2}, x_{i2}; \hat{\beta}_2). \tag{11}$$

3.3 Step 3: MWTP function inversion

At this point, the researcher is now equipped with two observations of implicit price and chosen level of q along the same MWTP curve for each individual (i.e. holding the individual's time-varying attributes fixed). In equilibrium, the implicit price will be equal to the MWTP function and this relationship allows me to write the problem of estimating linear MWTP curves as two equations with two unknowns (μ_{i0}, μ_{i1}):

$$\hat{P}^{q}_{i1} = \mu_{i0} + \mu_{i1} q_{i1} \tag{12}$$

$$\tilde{P}^{\tilde{q}}_{i2} = \mu_{i0} + \mu_{i1} \tilde{q}_{i2} \tag{13}$$

This system of two equations can be used to solve for the two unknowns (μ_{i0}, μ_{i1}) for each individual. Because individual heterogeneity is embodied in the preference parameters rather than in an additive regression error as is the case with Rosen's estimator I avoid the endogeneity problems described by Epple (1987) and Bartik (1987). This solution process amounts to finding the parameters of the MWTP function that "connect the dots" (CTD) for each individual.

$$\mu_{i1} = \frac{\tilde{P}^{\tilde{q}}_{i2} - \hat{P}^{q}_{i1}}{\tilde{q}_{i2} - q_{i1}} \tag{14}$$

$$\mu_{i0} = \hat{P}^{q}_{i1} - \mu_{i1} q_{i1} \tag{15}$$

In order to obtain identification, time variation in both implicit prices and demand for the amenity is needed. If prices are constant, one cannot detect how individuals respond to changing price schedules. Likewise, if $\tilde{q}_{i2} = q_{i1}$, the MWTP function parameters are also unidentified. To apply this method one must therefore have access to panel data on home purchases, prices and their buyers.

4 Data

In the application of this paper, I therefore exploit the rich Danish register datasets. As an example of the applicability of the CTD approach, I estimate the MWTP functions for violent crime. I describe the datasets in details below.

The data come from Statistics Denmark's confidential registers[5]. Overall, I use three types of data: housing data, individual demographic data and neighborhood data, all of which are observed on an annual basis. The housing and individual demographic data cover the period 1992-2015 while the neighborhood data cover a shorter time span such that I end up using data for 2008-2014. My analysis is conducted at the parish level. A parish is an administrative geographical unit which is used to assign individuals to a local church, where they have the right to having church ceremonies conducted. The boundaries of the parishes have mostly been stable for several hundreds years, though small changes have occurred over time due to building of new churches or joining of small parishes. I use the 2017 definition of 2,158 parishes in Denmark and keep that definition constant over time.

4.1 Individual demographic data

The datasets describing demographic information like home address, age, marital status, number of children (all from the population register BEF), education (from the register UDDA), home ownership (from the register EJER), and income (from the register IND) are merged based on the unique personal identifier PNR. While information from BEF is posted on January 1st of the year, EJER and UDDA are posted in the beginning of October each year and IND by the end of the year. To ensure that the observations are as close in time as possible, I merge BEF from year t together with UDDA, EJER and IND from $t - 1$.

From the population register I get information on all individuals living in Denmark for 1992-2015. I restrict attention to all home owners in the Copenhagen local labor market defined according to Statistics Denmark's definition from 2014, cf. Figure 3, from 2008-2014[6]. The address information I get from this register is an anonymized address for the street, number, floor and door such that apartment complexes consist of several unique addresses.

I only include home owners in the analysis because the decisions to own versus rent a house cannot be directly compared. Contrary to renters, home owners are making an actual investment in the property and may therefore consider how local amenities will evolve in the future. The hedonics literature typically ignores the dynamic component of the decision process (one notable exception is Bishop and Murphy (2011)), but I do focus on buyers to avoid confounding different objective functions). The information on home ownership comes from the ownership register EJER. Every property in Denmark has a record indicating which individuals own it and at what date they took over the ownership. I define home owners to be every person who owns more than 0% of an address in a calendar year.

[5] See Table A1 for an overview of the registers I use.

[6] Incorporating additional housing markets would complicate the model by requiring that I consider trade-offs in both labor and property markets, see Roback (1982).

Figure 3: Overview of municipalities in Denmark

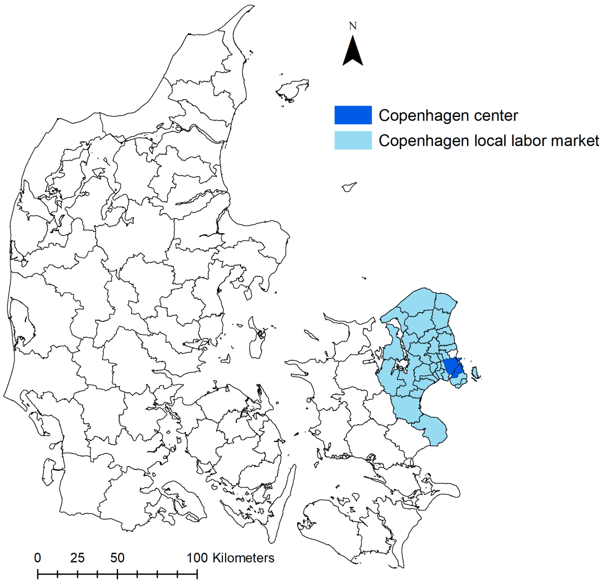

Note: Statistics Denmark's definition on local labor markets is based on municipalities, cf. Statistics Denmark (2016). Each municipality consists of several parishes.

4.2 Housing and transactions data

Data on housing characteristics come from the register BOL which holds a description of every housing unit in Denmark such as number of rooms, square meter living space, construction year, number of bathrooms, whether the building is historically preserved and if there is access to a kitchen. Sales prices come from the EJSA register and have been deflated to 2011 prices using the consumer price index. EJSA contains an observation for every housing unit sold including the transaction price, the type of sale (e.g. single-family house, commercial or farm property), number of square meters sold and the type of post-sale ownership (e.g. private, association, company or state). Lastly, I have data on valuations of all properties in Denmark from the register EJVK. These valuations are made by the tax authorities for property tax purposes every other year and consist of, among others, an assessment of the land value and the property value. EJER has a unique housing unit identifier which can also be found in the housing characteristics dataset BOL. I use this variable to merge EJSA and EJVK on to BOL which can then be merged on the personal dataset by using the address and PNR from BOL.

4.3 Neighborhood data

The neighborhood amenity data describe crime and school quality. From Statistics Denmark I access information on the number of victims by type of crime by year and parishes for 2005-2017. The types of crime are violent crime, sexual crime and property crime. More detailed crime types are available. For example, I can distinguish between burglary and theft in property crime. To avoid inclusion of types of crime that are unlikely to be reported or may not have anything do with the area itself (e.g. incestuous crime), I define violent crime to include serious violent crime, rape, crime against life and body, murder and attempted murder and violence against public authorities. The excluded groups are simple violence, threats and crime against personal freedom. I define property crime to include thefts and robberies and exclude blackmailing.

Despite having data on population, housing units and sales since 1992, I focus on the years 2008-2014 because I only have data on school districts[7] from Statistics Denmark's register SKOL for that sub-period. The school districts assign children to a primary and middle school (0th to 9th grade) based on their residential address. Hence, the school district data contain a link between all addresses in Denmark and a code for the school district that any home address belongs to in a given year. The school district boundaries can change over time and determine which public school parents are guaranteed to have their children accepted to. In theory there is free school choice implying parents are not forced to choose the local school, but to get their children into another school they must first apply and only if there are available seats will the parents' request be accepted. In the analysis, I exploit the school district boundaries to construct school fixed effects to account for the possibility that households sort based on school quality which then may influence house prices. These fixed effects also control for any other school district-specific amenities that potentially affect the house prices, but have been left out of the controls.

4.4 Sample selection

In the analysis, I only include sales that fulfill a number of criteria: the valuation of the property of the sale must exceed the value of the entire lot as the lot value is the value without any buildings and should therefore represent only a fraction of the property value. Further criteria are to only include sales that are not flagged as problematic by Statistics Denmark, the home is owned by a private individual or private housing cooperatives[8], the type of the sale belongs to one of the groups: single-family houses on private land, two-apartment houses or double houses on private land, three-apartment houses on private land, residential-only property with 4-8 apartments on private land, residential-only properties with 9 or more apartments on private land, mixed residential and business properties on private land excluding owner-occupied flats, developed farms, owner-occupied flats for residential use on private land, lots below 2000 square meter and other developed land. Excluded categories are: business-only properties, factories and warehouses, summer houses and other properties not belonging to any of the before-mentioned groups. Moreover, I only include addresses that have been sold once during the year, where only one household (family unit) lives, and where the parish code of the home is known as I use parishes to define neighborhoods. I also delete observations where the area sold is zero and for apartments where the area sold is above 500 square meters to avoid interpreting sales of whole apartment blocks as sales of single apartments. In general, I remove observations with sold area

[7]In Denmark each primary and middle school has an associated district, thus there is no distinction between school districts and catchment zones.

[8]Denmark has a tradition for housing cooperatives which are associations whose purpose is to buy, own and manage residential properties for the members of the association. Each member does not own his residence, but does own a share of the association and thereby the right to use one of its residences.

above the 99th percentile of the area distribution or if the number of rooms exceeds the 99th percentile of the rooms distribution[9].

Table 1 shows summary statistics for the property transactions I use, while Table 2 shows summary statistics for the buyers of properties in the Copenhagen local market during the period. In the estimation of the hedonic price function I use everyone with one or two purchases, while I only use those with two purchases when I identify individual-specific demand curves for violent crime.

To understand how much 1- and 2-purchase individuals differ in terms of characteristics, I compare the first part of Table 2 to the second and third part. Individuals who only buy one home in the period are on average as old as the two-purchase individuals are at the time of their second purchase. They are slightly more likely to be in a couple compared to both the first- and second-time buyers in the two-purchase individuals group. They also have a higher probability of having children compared to two-purchase individuals who buy their first home, but a lower probability compared to the time they buy their second home. In terms of education, the one-purchase home owners are more likely to be unskilled, and looking at household finances, the one-purchase individuals have a bit higher annual income, assets and debt than the two-purchase individuals at the time of the first purchase. However, this is reversed when comparing to statistics at the time of the second purchase. Lastly, the one-purchase individuals have a lower probability of moving job region in connection with the move and to buy a home in a big city compared to two-purchase individuals when they buy their first home. However, the job moving propensity pattern has flipped when doing the comparison at the time of the second purchase where the urbanization rate is also similar. Overall, the one-purchase individuals do not differ too much from the two-purchase individuals. Rather they represent an individual whose characteristics lie somewhat in between the characteristics of those who buy their first and second out of two homes in the period, and just seem to have postponed buying their first home until the time where they have more children and more settled with a partner. I therefore expect that estimating the model using data on the subsample of two-purchase individuals will provide estimates of MWTP that are relevant to the 1-purchase individuals as well.

Looking at Figure 4, there clearly is a significant amount of variation in the violent crime rates across homes in different parishes. The distributions do not move much across years, but mainly across space as depicted in Figure 5a. Here, the average number of victims of violent crime for 2008-2014 by parish in the Copenhagen local labor market are shown. The higher crime rates tend to be within a 10-15 km radius of the Copenhagen center. The same goes for property crime, cf. Figure 5b.

5 Results

With the rich Danish register data at hand, I can apply the CTD approach and estimate MWTP functions for violent crime rates. In this section I go over each of the steps involved in that procedure and explain the econometric specifications I use in practice.

[9] See Table A2 for an overview of the sample selection process

Table 1: Summary statistics of property transactions

	Mean	S.d	Median	N
Violent crime	10.18	16.51	6.00	58,920
Property crime	150.28	503.49	49.00	58,920
# sqm sold	475.51	453.00	347.00	58,920
$\mathbb{I}$[apartment]	0.37	0.48	0.00	58,920
$\mathbb{I}$[bath]	0.99	0.09	1.00	58,920
$\mathbb{I}$[preserved]	0.02	0.13	0.00	58,920
Build year	1956	33.02	1963	58,542
# rooms	4.02	1.39	4.00	58,920
Km to Copenhagen center	17.23	13.47	12.61	58,920
Inhabs. pr. km^2	3,748	5,456	1,720	58,920

Sample criteria: Only using one property observation within the household in the year. *Note:* $\mathbb{I}$ is the indicator function.

Table 2: Summary statistics of buyers at time of purchase by total number of purchases

	Mean	S.d	N
1 purchase			
Age	38.85	11.94	92,121
$\mathbb{I}$[couple]	0.83	0.38	92,121
$\mathbb{I}$[male]	0.50	0.50	92,121
$\mathbb{I}$[has children]	0.56	0.50	92,121
$\mathbb{I}$[has school age child]	0.23	0.42	92,121
Education			
Unskilled	0.04	0.18	92,121
High school	0.17	0.38	92,121
Vocational/Short Cycle Tertiary	0.25	0.43	92,121
Medium Cycle Tertiary	0.29	0.45	92,121
$\mathbb{I}$[divorce]	0.03	0.17	76,408
Household total inc. (10t DKK)	76.03	27.44	83,036
Household assets (10t DKK)	272.48	113.21	82,916
Household debt (10t DKK)	248.18	159.49	92,101
$\mathbb{I}$[new job municipality]	0.37	0.48	92,121
$\mathbb{I}$[live in big city]	0.54	0.50	92,121
2 purchases, 1st purchase			
Age	34.70	9.99	2,670
$\mathbb{I}$[couple]	0.77	0.42	2,670
$\mathbb{I}$[male]	0.51	0.50	2,670
$\mathbb{I}$[has children]	0.45	0.50	2,670
$\mathbb{I}$[has school age child]	0.16	0.36	2,670
Education			
Unskilled	0.03	0.17	2,670
High school	0.21	0.41	2,670
Vocational/Short Cycle Tertiary Tertiary	0.23	0.42	2,670
Medium Cycle Tertiary	0.28	0.45	2,670
Long Cycle Tertiary	0.24	0.43	2,670
Household total inc. (10t DKK)	71.45	26.83	2,438
Household assets (10t DKK)	267.42	113.41	2,375
Household debt (10t DKK)	242.71	202.40	2,670
$\mathbb{I}$[new job municipality]	0.42	0.49	2,670
$\mathbb{I}$[live in big city]	0.69	0.46	2,670
2 purchases, 2nd purchase			
Age	38.07	9.73	2,670
$\mathbb{I}$[couple]	0.80	0.40	2,670
$\mathbb{I}$[male]	0.51	0.50	2,670
$\mathbb{I}$[has children]	0.65	0.48	2,670
$\mathbb{I}$[has school age child]	0.23	0.42	2,670
Education			
Unskilled	0.02	0.15	2,670
High school	0.13	0.34	2,670
Vocational/Short Cycle Tertiary	0.25	0.43	2,670
Medium Cycle Tertiary	0.29	0.45	2,670
Long Cycle Tertiary	0.30	0.46	2,670
Household total inc. (10t DKK)	79.19	28.71	2,447
Household assets (10t DKK)	282.13	114.90	2,461
Household debt (10t DKK)	283.56	163.53	2,670
$\mathbb{I}$[new job municipality]	0.34	0.47	2,670
$\mathbb{I}$[live in big city]	0.54	0.50	2,670

Note: $\mathbb{I}$ is the indicator function. $\mathbb{I}$[new job municipality]$= 1$ if either or both of the household members gets a job in t in another municipality than where they had a job in $t-1$. Monetary terms deflated by 2011 consumer price index.

Figure 4: Probability density function of number of victims of violent crime per 1,000 people

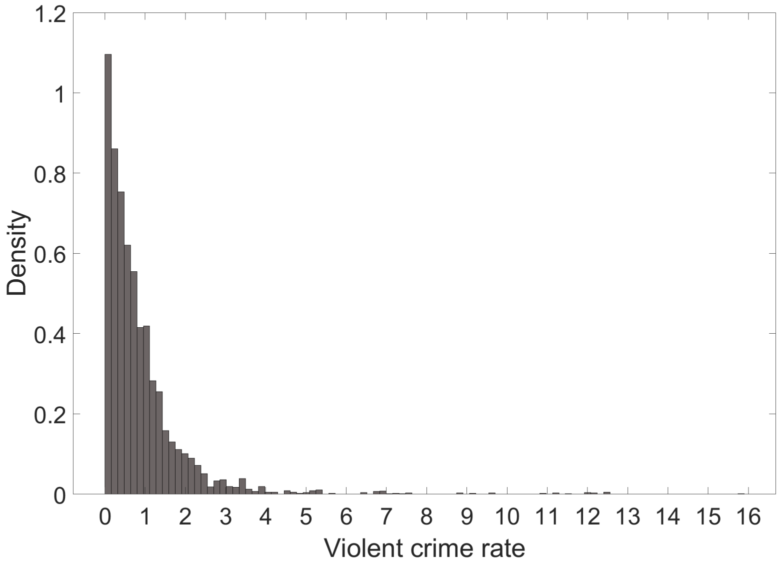

Note: Violent crime rate is defined as the number of victims of violent crime per 1,000 people.

Figure 5: Average number of victims of crime 2008-2014 by parish in Copenhagen local labor market

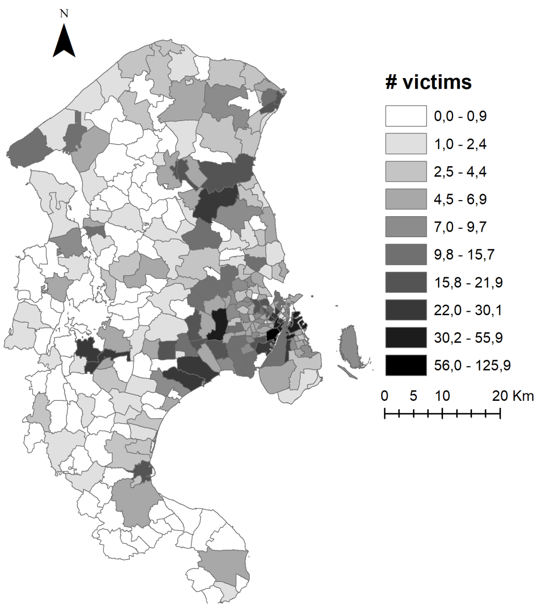

(a) Violent

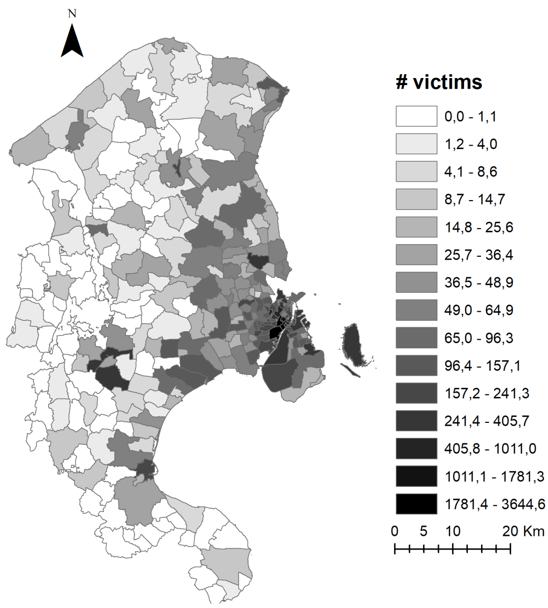

(b) Property

5.1 Step 1: Hedonic gradient

Following Bishop and Timmins (2019) I model the hedonic price function semi-parametrically according to

$$P_{it}(q_{it}, x_{it}; \beta_t) = x'_{it}\beta^x_t + \Lambda(q_{it}; \beta_t) + \epsilon_{it}, \tag{16}$$

where q_{it} is a scalar describing the amenity of interest (i.e. violent crime pr. 1,000 inhabitants), $\Lambda(.)$ a flexible function of q_{it}, ϵ_{it} a regression error and x_{it} a vector of other housing or neighborhood attributes. I control for the latter in order to obtain an estimate of the causal effect of q_{it} on price. It holds quadratic functions of the property crime rate, square meters sold and number of rooms as well as dummy variables for whether the housing unit sold has a bathroom, has a kitchen, if it is an apartment and lastly a set of school district fixed effects. Kuminoff et al. (2010) do indeed show that inclusion of spatial fixed effects is the preferred way to control for unoberved neighborhood amenities.

To minimize the risk of bias from functional form, I choose to model P_{it} as a flexible function of q_{it} using the semi-parametric method in Robinson (1988). Hence, I first obtain a predicted local conditional mean of x_{it} and P_{it}, respectively, using local polynomial regression on q_{it}. I then construct residuals for each of these variables by subtracting their predicted local means from x_{it} and P_{it}. Express these residuals by a dot: $\dot{x}_{it}$ and $\dot{P}_{it}$. I next finalize the first stage of Robinson's procedure by obtaining an estimate of β^x_t by regressing $\dot{P}_{it}$ on $\dot{x}_{it}$ using OLS. With β^x_t at hand, I can proceed to the second stage of Robinson's method and move $x'_{it}\beta^x_t$ to the left side of Equation 16 in order to regress $P_{it} - x'_{it}\hat{\beta}^x_t$ on q_{it} using local polynomial regression.

I aim at allowing for maximum flexibility while making sure the resulting price functions are consistent with utility maximization. Consequently, I perform a careful specification search over the global and adaptive bandwidth parameters as well as the local polynomial degree. I smooth the hedonic price function up until the point where all points exhibit a negative gradient. During this search I note that a local polynomial degree of more than 1 results in wiggly gradients which are inconsistent with utility maximization. Thus, I estimate the price function using a linear local polynomial function of the violent crime rate. I end up concluding that a global bandwidth of 2.5 times the standard deviation of violent crime and an adaptive bandwidth of 0 achieves my goal of negative gradients across the crime distribution with as little smoothing as possible[10]. This is in line with Bishop and Timmins (2019) who use a bandwidth equal to 2.15 times the standard deviation of the violent crime rate.

The results from this first-step estimation are illustrated in Figure 6. It shows the price function (panel a) and the gradient (panel b) as a function of violent crime. Figure B1 shows similar results for the gradients including 99% bootstrapped confidence intervals. As expected, the hedonic price function is positive and slopes downwards. That is, all gradients are negative as depicted in Figure 6b. The gradients do show variability over time and thereby provide the variation necessary to identify the MWTP curves. However, these gradients are also non-decreasing in violent crime and thereby illustrate the problem of using gradients directly as a

[10] I also apply the same bandwidths and polynomial degree in the first stage of Robinson's method.

Figure 6: Results of 1st Stage by Year, $P_t(Z_{i,t})$

(a) Hedonic price function (b) Hedonic gradient

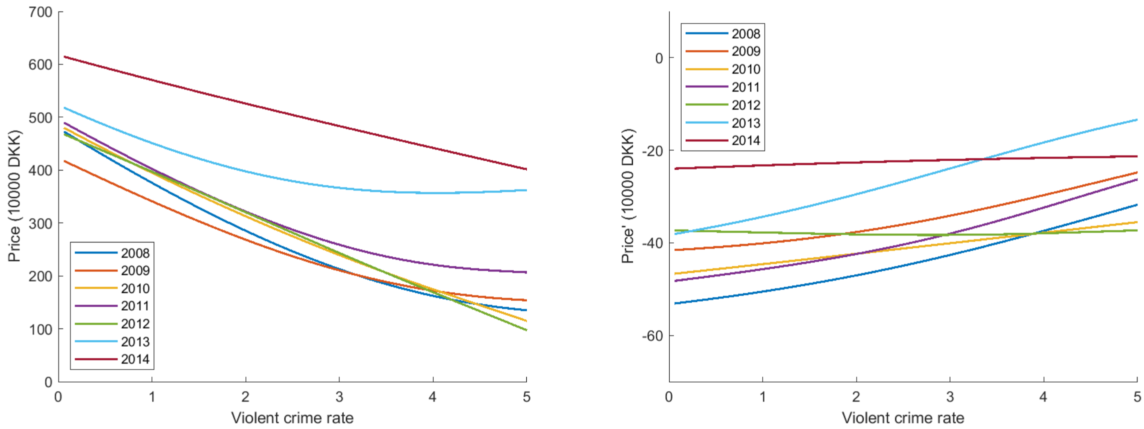

Note: The violent crime rate is measured as number of victims of violent crime per 1,000 people.

measure of MWTP as they would indicate that violent crime is a good. This will be demonstrated in more detail using the Rosen method below.

There are a number of other papers in the literature that study how crime affects house prices, but the most comparable in terms of the specification for the first stage and the definition of the violent crime rate is Bishop and Timmins (2019). They estimate gradients of the rental equilivalent of housing price (5% of the total house price in 2000 USD) with respect to the number of victims of violent crime per 100,000 people. They find a hedonic gradient in the range -40 to 0. Converting these estimates into the corresponding gradient of house prices in 10,000 DKK with respect to crime per 1,000 individuals gives estimates in the range 0 to -70[11] which is completely in line with the estimates in Figure 6b.

5.2 Step 2: Segmentation equations

Next, I estimate the segmentation equation in Equation 9 and exploit the panel dimension in the Danish census to include individual-specific fixed effects. This allows me to capture that individuals may have time-constant unobserved traits that affect their demand for less crime (demand for more safety). For instance, individuals who have grown up with very little crime might have higher demand for safety than other individuals. Mapping out all potential reasons why individuals may be more or less averse to crime is unrealistic and controlling for individual fixed effects solves this issue for time-constant unobservables. To add flexibility to the counterfactual demand predictions over time, I include calendar year effects. I performed a specification search by including a number of time-varying individual characteristics observed in the data and also attempted interactions between these and the year effects and also interactions

[11] Denote estimates using their definition as $gP_{USD,crime100,000}$ and my definition $gP_{DKK,1,000}$. The USD-DKK exchange rate in 2000 was approximately 8.7. The conversion is computed as $gP_{DKK,1,000} = gP_{USD,crime100,000}/0.05 \cdot 8.7 \cdot (100,000/1,000)/DKK_{scale}$ where I use $DKK_{scale} = 10,000$.

between the time-varying variables themselves (not shown). My preferred specification is column (1) in Table 3. It controls for a continuous, linear year effect and child dummies. Including year fixed effects instead as in column (5) resulted in insignificant coefficients. This is not too surprising given that I use data on individuals buying at least two homes during 2008-2014 and conditioning on year dummies leaves rather few observations in each year. Hence, the year effects will be harder to estimate precisely. Allowing for a non-linear effect of calendar year as in column (2), changes in total household income as in column (3) or changes in martial status through divorce as in column (4) do not improve the specification significantly. The main driver of changes in demand for less crime is therefore changes in family status through children and the general changes over time in the attitude towards living in neighborhoods with crime.

The time-varying amount of children does have an economically significant effect on the demand for crime. The baseline demand for crime for someone in 2008 with 0 children is 0.818[12]. Adding an effect of children between -0.45 and -0.35 changes the demand by a relatively large amount. The same can be concluded for the individual fixed effects. The distribution of these estimates is shown in Figure 7. The majority of the probability mass is between -1 and 1, hence increasing or decreasing the baseline level of 0.818 by more than 100%. These findings underline the value of having access to panel register data where it is possible to account for both individual fixed unobservables and changes in background characteristics that significantly affect the demand for the (dis)amenity of interest.

To predict the individual's chosen quantity of violent crime at the time of the second purchase were she to have the attributes she had at the time of the first purchase, I first adjust the individual back to those attributes and then use the estimates from Table 3 for the prediction including the estimate of her indvidual fixed effect.

[12] Computed by $-0.043 \cdot 2008 + 87.162$.

Figure 7: Distribution of fixed effects in demand for violent crime

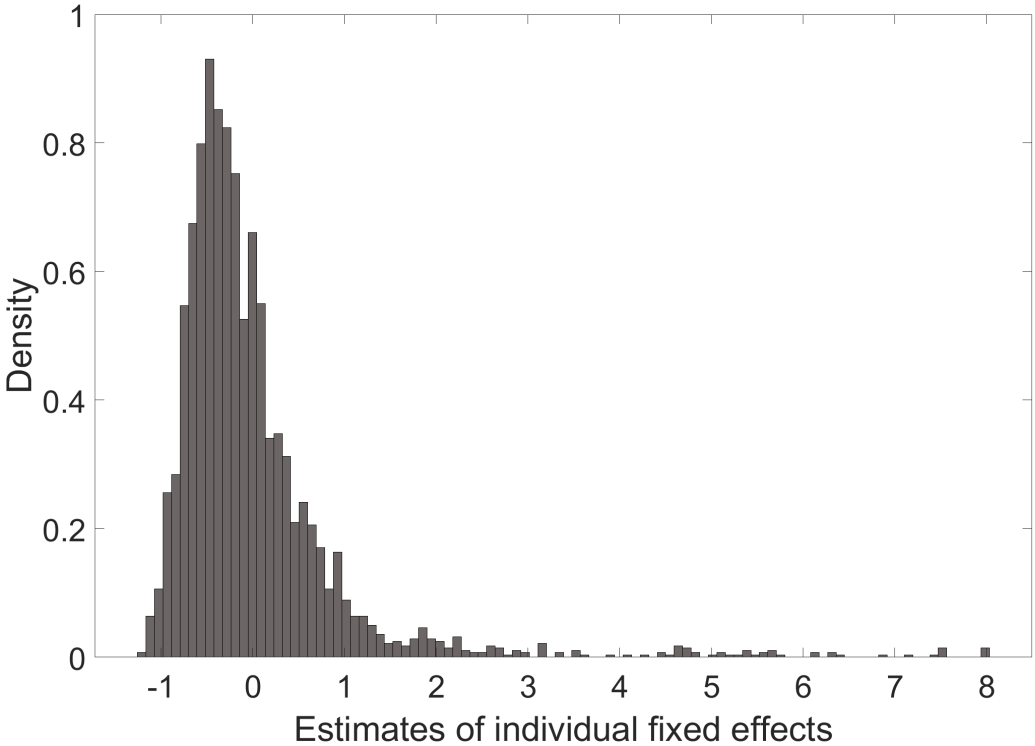

Note: The violent crime rate is measured as number of victims of violent crime per 1,000 people. Estimates come from specification (1) in Table 3.

Table 3: Segmentation equation for violent crime rate with individual fixed effects

	(1)	(2)	(3)	(4)	(5)
Year	-0.043***	-17.737	-0.040***	-0.041***	
	(0.01)	(23.10)	(0.01)	(0.01)	
Year^2		0.004			
		(0.01)			
Number of children (ref. 0)					
1 child	-0.352***	-0.350***	-0.336***	-0.351***	-0.348***
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
2 children	-0.388***	-0.387***	-0.366***	-0.393***	-0.385***
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)
3+ children	-0.448***	-0.447***	-0.420***	-0.454***	-0.452***
	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
Household income (10,000 DKK)			-0.001		
			(0.00)		
Household income (10,000 DKK)^2			0.000		
			(0.00)		
$\mathbb{I}[Divorce]$				0.085	
				(0.09)	
$\mathbb{I}[Year = 2009]$					-0.108
					(0.08)
$\mathbb{I}[Year = 2010]$					-0.200***
					(0.07)
$\mathbb{I}[Year = 2011]$					-0.075
					(0.08)
$\mathbb{I}[Year = 2012]$					-0.301***
					(0.08)
$\mathbb{I}[Year = 2013]$					-0.222***
					(0.08)
$\mathbb{I}[Year = 2014]$					-0.301***
					(0.07)
Constant	87.162***	17,880.2	82.509***	83.524***	1.313***
	(18.83)	(23228.41)	(18.68)	(18.92)	(0.05)
N	6,167	6,167	6,167	6,167	6,167

Sample criteria: Individuals buying a property at least twice during 2008-2014. *Note:* Estimated by OLS with individual fixed effects. Standard errors in parentheses clustered at the individual level. Violent crime rate measured as number of victims per 1,000 people. $\mathbb{I}$ is the indicator function. Average household income is 81.5 (10,000 DKK).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.3 Step 3: MWTP function inversion

Having obtained the predicted demand for violent crime at the second purchase, I derive an estimate of the implicit price the individual would have had to pay for that level of violent crime in the year when the second purchase took place. To do this, I use nearest neighbor interpolation of the gradient function[13]. With both implicit prices and demand available for each period, I can now evaluate Equation 14 and Equation 15 and thereby identify the individual's MWTP function for reductions in violent crime.

5.4 Interpretation of results

In order to summarize the results of the many heterogeneous MWTP functions that I recover, I illustrate the distribution of μ_0, μ_1 and MWTP in Figure 8. μ_0 is distributed across the range -50 to 250, i.e. a large number of indviduals have positive μ_0. This may seem counterintuitive if interpreting μ_0 as the MWTP for reductions in crime at a crime level of 0, but is in line with results in the literature, e.g. Bishop and Timmins (2019). In any case, the focus of the paper is on the slopes of the MWTP function, not the levels themselves and all individual MWTP curves are indeed consistent with utility maximization in the sense that they comply with the single-crossing property. As expected from the theory, μ_1 is negative on average for everyone. The MWTP for reductions in crime (more safety) is positive and in the range of 200,000-550,000 DKK. That is, safety is considered a good. The distribution of MWTP shows several peaks reflecting the large amount of heterogeneity in MWTP for reductions in crime across individuals. Regressing the negative MWTP on children, income, and year I find that individuals with more children have a higher MWTP for reductions in crime, all else equal, cf. Table 4. This is consistent with them demanding less crime as found in the segmentation equation. These summarizing results also show that over time, the MWTP for reductions in crime has generally decreased. Again, this is in line with the trend that individuals tend to reside in parishes with less crime over time and therefore value further reductions in crime by less.

[13] Results are not sensitive to the choice of interpolation method.

Figure 8: CTD: Results of Inversion for μ_0, μ_1 and Negative MWTP

(a) μ_0 (b) μ_1 (c) -MWTP for violent crime

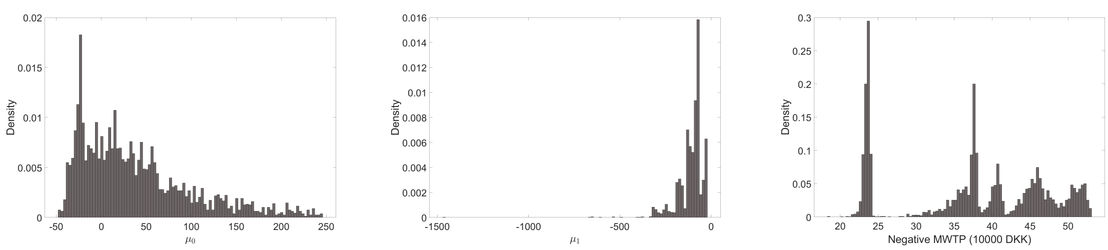

Note: Removing individuals with the 5% most extreme estimates of MWTP. Violent crime measured as number of victims of violent crime per 1,000 people.

Table 4: CTD: OLS of negative MWTP

Variable	1 kid	2 kids	3+ kids	2009	2010	2011	2012	2013	2014	Cons.
Coef.	0.196	0.559***	0.869***	-10.221***	-5.712***	-5.459***	-12.706***	-15.799***	-26.712***	49.803***
S.e.	(0.12)	(0.12)	(0.14)	(0.14)	(0.16)	(0.27)	(0.13)	(0.15)	(0.10)	(0.11)

Note: $N = 5,608$. Violent crime rate measured as number of victims per 1,000 people.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.4.1 Analysis using Rosen (1974)

I use the estimates of the hedonic price function from Figure 6 combined with information about individual homeowners using the theory described in Section 2. Estimates reported in Table 5 come from Rosen's second-stage regression, using the implicit price of violent crime derived from the hedonic gradient as the dependent variable. Hedonic gradients for each individual have been computed by interpolating the hedonic gradient function as in section 5.0.2. I use the same set of observations as in my second stage segmentation equation in Table 3 and also control for individual-specific fixed effects to enable a reasonable comparison between the two methods. The endogeneity concern raised by Epple (1987) arises because the individual's choice of violent crime is determined by the same unobserved preference shock that determines the implicit price of violent crime because of the process of sorting along a non-linear budget constraint. The concern in that setting is that the coefficient on violent crime will have a positive bias. That is indeed what I find. That bias is severe enough that in every specification, the coefficient on violent crime (i.e., the slope of the MWTP function) is positive and statistically significant. In the following section, I demonstrate how the counterintuitive slope affects my welfare measure associated with non-marginal changes in violent crime.

Table 5: Rosen 2nd stage: OLS of MWTP with individual fixed effects

	(1)	(2)	(3)	(4)
Violent crime rate	1.671***	2.174***	1.854***	2.103***
	(0.13)	(0.18)	(0.12)	(0.19)
Number of children (ref. 0)				
1 child	-0.848**			-0.892***
	(0.34)			(0.34)
2 children	-0.744**			-0.662*
	(0.37)			(0.37)
3+ children	-1.956***			-1.929***
	(0.62)			(0.62)
1 child × Violent crime rate	0.516*			0.583*
	(0.31)			(0.31)
2 children × Violent crime rate	0.543**			0.598**
	(0.27)			(0.27)
3+ children × Violent crime rate	0.664			0.896*
	(0.53)			(0.53)
Year	3.754***	3.740***	3.694***	3.742***
	(0.03)	(0.03)	(0.03)	(0.04)
Household income (10,000 DKK)		-0.001		0.000
		(0.00)		(0.00)
Violent crime rate × Household income (10,000 DKK)		-0.004**		-0.005**
		(0.00)		(0.00)
$\mathbb{I}[divorce]$			-1.374***	-1.308***
			(0.43)	(0.44)
$\mathbb{I}[divorce]$ × Violent crime rate			-0.111	-0.160
			(0.33)	(0.34)
Constant	- 7,589.319***	- 7,561.693***	- 7,469.329***	- 7,564.777***
	(69.18)	(61.84)	(60.39)	(70.41)
N	6,167	6,167	6,167	6,167

Sample criteria: Individuals buying at least two properties during 2008-2014. *Note:* $\mathbb{I}$ is the indicator function. The dependent variable is the estimated implicit price of violent crime using estimates from Figure 6. Standard errors in parentheses clustered at the individual level. Violent crime rate measured as number of victims of violent crime per 1,000 people. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.5 Welfare analysis

In this section I demonstrate the consequences of mis-measuring the slope of the MWTP function by comparing the value of a large change in violent crime derived using Rosen's two-stage approach to that using my CTD procedure. In practice, I calculate the willingness to pay (WTP) to avoid a 30 percent increase using CTD and Rosen's method.

The WTP is the area between the MWTP curve and the horizontal axis between the current and new level of crime. Let q_0 denote the current level of crime and q_{high} the level of crime after the 30 percent increase, i.e. $q_{high} = q_0 \cdot 1.3$. Figure 9 illustrates the concept: the WTP to avoid an increase in violent crime from q_0 to q_{high} using the CTD procedure corresponds to area (4)+(5)+(6), while it corresponds to area (4)+(5) if I just assumed a horizontal MWTP curve[14] and area (4) if I used Rosen's method. Likewise, if I considered a reduction in violent crime from q_0 to q_{low}, I would get a WTP of area (1) if I used the CTD approach, area (1)+(2) if I used the

[14] That is the implication of Bajari and Benkard (2005) unless the researcher assumes an explicit function form for preferences that ensures downwards-sloping demand curves, e.g. a Cobb-Douglas utility function.

horizontal MWTP function and area (1) + (2) + (3) if I relied on Rosen's method. I.e. because the Rosen model predicts upward-sloping MWTP functions for violent crime, the MWTP for avoiding further crime increases drops as the level of crime increases, meaning that the Rosen model predicts a smaller cost from a violent crime increase than does the CTD approach. The opposite argument can be used to show that the Rosen model predicts a larger benefit from a violent crime reduction.

Figure 9: Example: computing WTP using different methods

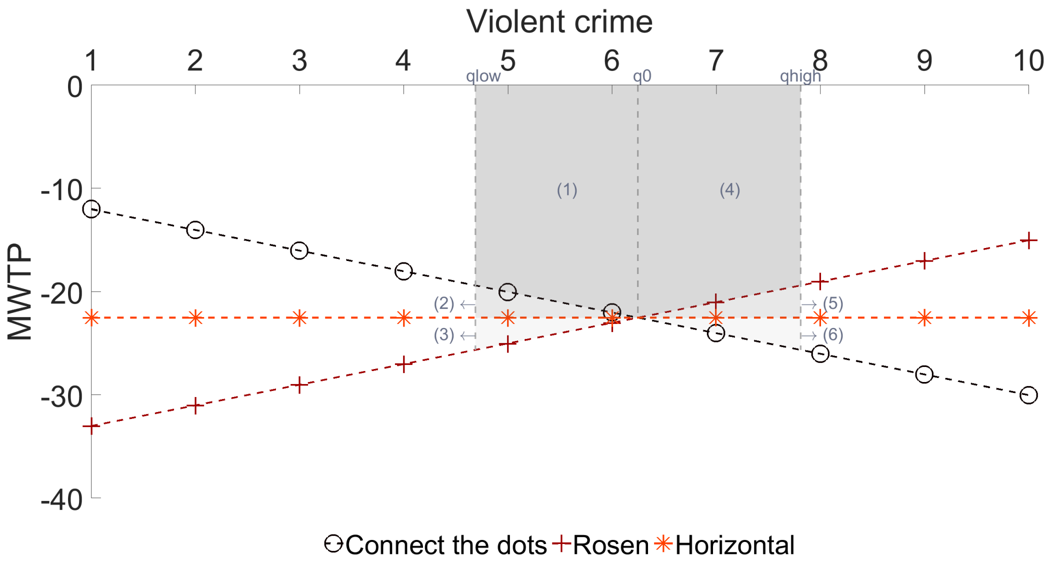

Algebraically, the MWTP function for individual i using the CTD approach is given as

$$MWTP_{it}^{CTD} = \mu_{i0} + \mu_{i1} q_{it} \tag{17}$$

The WTP to avoid an increase in violent crime for i is then calculated for each individual by

$$\begin{aligned} WTP_{it}^{CTD} &= -\int_{q_{high}}^{q_0} (\mu_{i0} + \mu_{i1} q) dq \\ &= -(\mu_{i0} \cdot (q_{it,high} - q_{it,0}) + 0.5 \cdot \mu_{i1}(q_{it,high}^2 - q_{it,0}^2)). \end{aligned} \tag{18}$$

For Rosen's method I use the estimate from the Rosen second stage where the MWTP, i.e. the estimated implicit price, has been regressed on the level of crime:

$$MWTP_{it}^{R} = \alpha_0 + \alpha_1 \cdot q_{it} + \epsilon_{it}, \tag{19}$$

where α_0 controls for everything that does not vary with q_{it} and α_1 controls for everything that does according to the regressions in Table 5. ϵ_{it} is a regression error. In the comparison between CTD and Rosen, I use specification (1) from Equation 19 since household income in specification (2) enters insignificantly and so does the divorce indicator in specification (3). Specification (1) is also similar to the specification used in the segmentation equation for CTD.

To get the WTP for Rosen's method, I integrate Equation 19 between q and q_{high}:

$$WTP_{it}^{R} = -(\alpha_0 \cdot (q_{it,high} - q_{it,0}) + 0.5 \cdot \alpha_1(q_{it,high}^2 - q_{it,0}^2)). \tag{20}$$

Figure 10: Bias of Rosen's negative WTP for a 30% increase in violent crime

(a) 10,000 DKK

(b) % difference

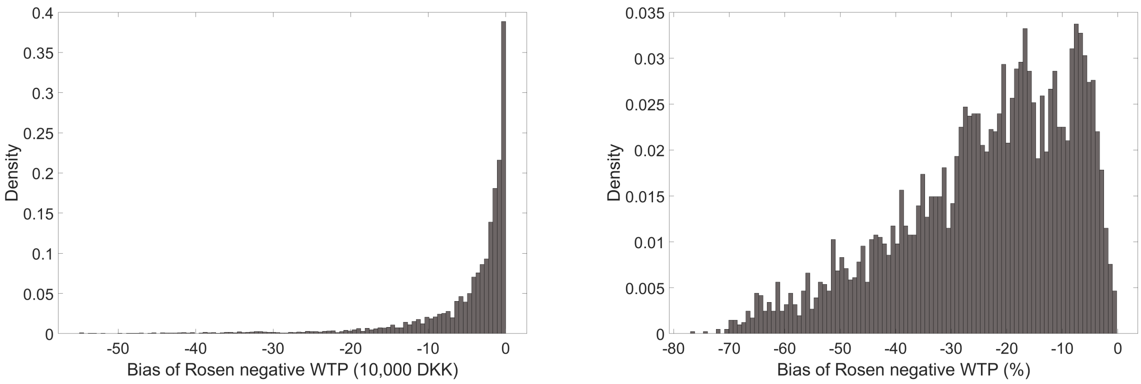

Note: The violent crime rate is measured as number of victims of violent crime per 1,000 people.

There is a lot of heterogeneity in the WTP across individuals as evidenced above and this is also replicated in the bias. Figure 10a shows the distribution of the bias in the negative WTP for violent crime in 10,000 DKK. As expected, I find the bias to be negative, i.e. Rosen underestimates the cost of crime as hypothesized in Figure 9. The bias is very economically significant and varies from just below 0 to more than -100,000 DKK. These magnitudes correspond to a bias of up to -70% according to Figure 10b and on average -24.4%.

Exploring how the bias varies with crime level, I plot violent crime against the magnitude of the bias in Figure 11. There is a clear negative correlation hence providing evidence that the bias is even larger for individuals who are living in high-crime areas. In the most crime-intensive areas, the bias is up to -500,000 DKK. This raises an even bigger concern from a policy-perspective as Rosen's model would underestimate the costs of letting crime increase in these areas even more than it does for more safe areas. Rosen's method would therefore tend to channel resources, that hinder crime from increasing, away from high-crime areas to a larger extent than it does for the safer areas. The distriubtion of such resources aross parishes would therefore be non-optimal and the welfare burden of using wrong estimates of the WTP for safety would fall more heavily on inhabitants in parishes with more crime.

6 Conclusion

For many years, the hedonics literature has struggled with how to recover preferences underlying the choices observed in the housing market. Accurately recovering these preferences is necessary for measuring the value of non-marginal changes in (dis)amenities. Because of the econometric problems described above, simple "first-stage" (in the parlance of Rosen (1974)) techniques have been used instead, but these methods only provide valid approximations for marginal changes in amenities of local public goods. Policy-relevant changes tend to be non-marginal. Over the last two decades, a number of techniques have been developed to address this problem. This paper

Figure 11: Bias of Rosen's negative WTP for a 30% increase in violent crime

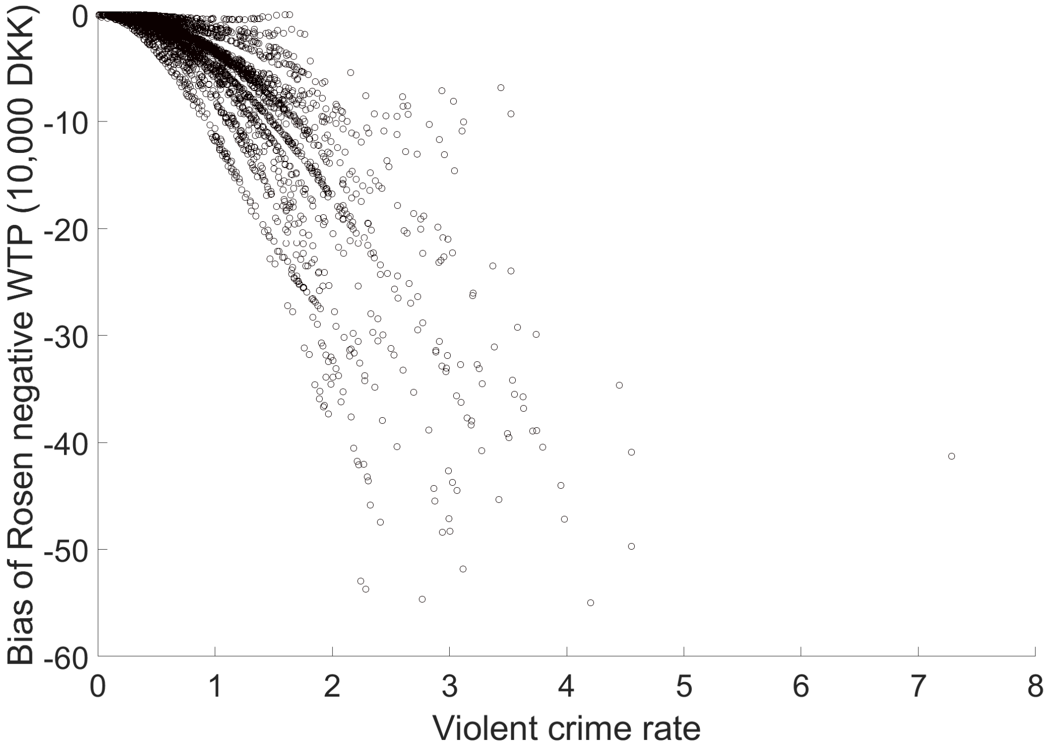

Note: The violent crime rate is measured as number of victims of violent crime per 1,000 people.

contributes to that literature, extending the analysis in Bajari and Benkard (2005) to allow individual MWTP functions to have both heterogeneous intercepts *and* slopes. In so doing, I extend the method developed by Bishop and Timmins (2018) to incorporate rich information about time-varying individual attributes making more practical the use of repeat buyer data to address the non-marginal valuation question. I implement that method using detailed data from the Danish census and find considerable heterogeneity in MWTP across the distribution of Danish households. Applying my model to valuing large increases in violent crime and comparing it to Rosen's procedure, I find that the Rosen model does indeed lead to biased estimates of the MWTP function and that these biases (on the order of 24 percent on average and up to 70 percent) are significant, increase in the level of crime and very policy-relevant.

7 Appendices

Appendix A Overview of data

Table A1: Overview of confidential registers from Statistics Denmark

Register	Description	Availability	Update
BEF	Main register listing all individuals with official address in Denmark. Information on basic information like social security number (SSN), age, home address, marital status, spouse's SSN, country of origin, and children and their SSN, parents' SSN and gender.	1992-current	January 1st
UDDA	Education register with information on highest obtained education of the individual including code for the educational institution and detailed fields and levels of study.	1992-current	October 1st
IND	Income register with information on annual total income, total wage income, assets, debt, public transfers and tax payments.	1992-current.	December 31st
EJER	Property ownership register with information on the SSN of owner, property identification code, type of ownership (e.g. public or private), ownership share of property and start date of ownership.	1992-current	October 1st
BOL	Property census register with information on every property unit in Denmark, e.g. number of rooms, living space , type of property, address and construction year. The register is based on The Central Register of Buildings and Dwellings (BBR) which is used for property assesments.	1992-current	January 1st
EJSA	Property transactions register with information on transactions of all real properties in Denmark such as the transactions prices, type of sale, land value, square meters sold and property identification code.	1992-current	January 1st
EJVK	Property assesment register based on BBR. As a general rule, the tax authorities assess the value of all owneroccupied dwellings in uneven years and other dwellings in even years.	1992-current	October 1st
KROF	Register of reports of victims of criminal offense by type of offense with information on e.g. address of the crime scene, victim's SSN and gender. We have not used individual-level data on crimes but instead got Statistics Denmark to deliver a dataset holding the number of victims by type of crime, parish and year based on data in KROF.	2001-current, but detailed address of crime scene incl. parish only 2005-current	January 1st
SKOL	Register of school districts with information on the home addresses that belong to the district. Data is based on CPR Vejregister which is a complete registry of all roads in Denmark including certain distric divisions such as school district. Municipalities decide the school districts and report these to CPR Vejregister.	2007-current	January 1st

Table A2: Overview of sample selection process

Selection criteria	N
Main dataset (all adult individuals in Denmark 2008-2014, sales and no sales)	29,460,516
Non-missing sales price (i.e. potential sales observation)	1,265,418
Property value > Lot value and private sale	664,928
Rooms < 99th percentile	662,776
Area sold < 99th percentile	656,439
Sale marked as OK by Statistics Denmark	656,439
Private owner post sale	591,492
Sales type either single-family houses on private land, two-apartment houses or double houses on private land, three-apartment houses on private land, residential- only property with 4-8 apartments on private land, residential-only properties with 9 or more apartments on private land, mixed residential and business properties on private land excluding owner-occupied flats, developed farms, owner-occupied flats for residential use on private land, lots below 2000 square meter and other developed land	582,920
Max number of distinct sales of the property: 1	580,395
Max number of sales of the property on the same date: 1	518,207
Property sold on open market terms	513,655
Property is sold in the current year	513,655
Max number of households (family units) on the address: 1	443,179
Parish of the address of the sale is known	443,179
Property sold to household who lives on address	314,199
Property was bought in Copenhagen local labor market	115,882
Individual's education known	99,779

Appendix B Bootstrap results of 1st-stage estimates

Figure B1: Hedonic Gradient with Bootstrap Confidence Intervals by Year, $P_t(Z_{i,t})$

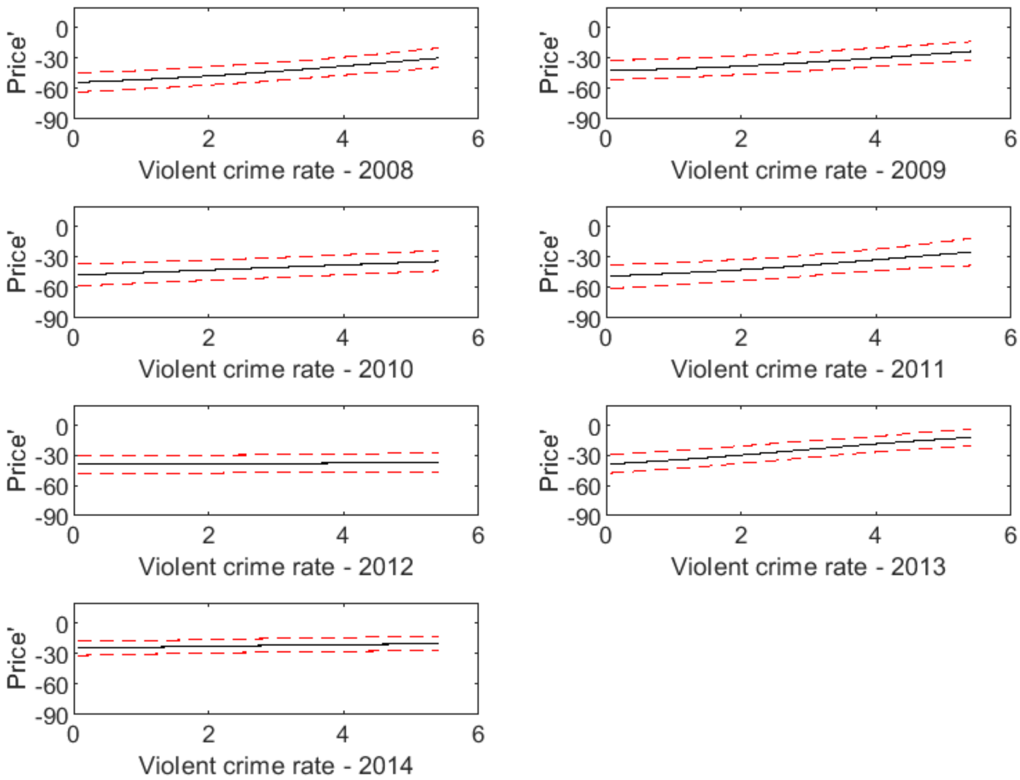

Note: Violent crime rate is measured as number of victims of violent crime per 1,000 people. Dashed lines are 99% confidence intervals.

8 Acknowledgments

I thank Innovation Fund Denmark (grant number: 5156-00001B) for financing access to data from Statistics Denmark and Rebecca Udsen for research assistance on checking crime data for flaws. A special thanks to Christopher Timmins, Kelly Bishop and Spencer Banzhaf for discussions of ideas for this project and to Ismir Mulalic for providing comments.

Bibliography

Bajari, P. and C. L. Benkard (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of political economy 113*(6), 1239–1276.

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

Bishop, K. C. and A. D. Murphy (2011). Estimating the willingness to pay to avoid violent crime: A dynamic approach. *American Economic Review 101*(3), 625–29.

Bishop, K. C. and C. Timmins (2018). Using panel data to easily estimate hedonic demand functions. *Journal of the Association of Environmental and Resource Economists 5*(3), 517–543.

Bishop, K. C. and C. Timmins (2019). Estimating the marginal willingness to pay function without instrumental variables. *Journal of Urban Economics 109*, 66–83.

Brown, J. N. and H. S. Rosen (1982). On the estimation of structural hedonic price models. *Econometrica 50*(3), 765–768.

Ekeland, I., J. J. Heckman, and L. Nesheim (2004). Identification and estimation of hedonic models. *Journal of political economy 112*(S1), S60–S109.

Epple, D. (1987). Hedonic prices and implicit markets: estimating demand and supply functions for differentiated products. *Journal of political economy 95*(1), 59–80.

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.

Mendelsohn, R. (1985). Identifying structural equations with single market data. *The Review of Economics and Statistics 67*(3), 525–529.

Roback, J. (1982). Wages, rents, and the quality of life. *Journal of political Economy 90*(6), 1257–1278.

Robinson, P. M. (1988). Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, 931–954.

Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy 82*(1), 34–55.

Statistics Denmark (2016). Pendlingsområder - metode. Technical report.

Willig, R. D. (1976). Consumer's surplus without apology. *The American Economic Review 66*(4), 589–597.