

NBER WORKING PAPER SERIES

PANEL DATA HEDONICS:
ROSEN'S FIRST STAGE AND DIFFERENCE-IN-DIFFERENCES AS "SUFFICIENT STATISTICS"

H. Spencer Banzhaf

Working Paper 21485
<http://www.nber.org/papers/w21485>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2015

I thank David Albouy, Olivier Beaumais, Jared Carbone, Kerry Smith, Reed Walker and Jeffrey Zabel for helpful comments and suggestions on previous drafts. I am especially grateful to Paul Ferraro and Nick Kuminoff for very detailed comments and discussions. I also thank seminar participants at the Univ. of Illinois, Univ. of South Carolina, and Yale Univ. and participants at an Urban Economics Association session at the 2013 NARSC meetings (Atlanta, GA), the 2013 National Tax Association meetings (Tampa, FL), the 2014 Workshop on Nonmarket Valuation (Aix-en-Provence, France), and the 2015 NBER Summer Institute. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by H. Spencer Banzhaf. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Panel Data Hedonics: Rosen's First Stage and Difference-in-Differences as "Sufficient Statistics"
H. Spencer Banzhaf
NBER Working Paper No. 21485
August 2015
JEL No. D46,D61,H4,Q51,R3

ABSTRACT

For decades, economists have used the hedonic model to estimate demands for the implicit characteristics of differentiated commodities. The traditional cross-sectional approach can recover marginal willingness to pay for characteristics, but has faltered over a difficult endogeneity problem for non-marginal welfare measures. I show that when marginal prices can be reliably estimated, and when panel data on household demands is available, one can construct a second-order approximation to non-marginal welfare measures using only the first-stage marginal prices. Under a single-crossing restriction, the approach remains valid for repeated cross sections of product prices.

More recently, economists have questioned the assumptions under which one can identify these cross-sectional hedonic price functions, raising the possibility of unobservables that are correlated with the characteristic of interest. To overcome this problem, they have introduced difference-in-differences econometric models to identify capitalization effects. Unfortunately, the interpretation of these effects has not been clearly perceived in the literature. I additionally show these capitalization effects are the "average direct unmediated effect" on prices of a change in characteristics, which can be interpreted as a movement along the ex post hedonic price function. This effect is a lower bound on Hicksian equivalent surplus.

H. Spencer Banzhaf
Department of Economics
Andrew Young School of Policy Studies
Georgia State University
P.O. Box 3992
Atlanta, GA 30302
and NBER
hsbanzhaf@gsu.edu

Panel Data Hedonics: Rosen's First Stage and Difference-in-Differences as "Sufficient Statistics"

1. Introduction

For decades, the hedonic model has been the starting point for understanding people's values for differentiated products. Its applications include willingness to pay in higher housing prices for local public goods and spatial amenities, compensating wage differentials for attributes like job safety, pricing of quality-differentiated consumer products like computers and cars, and quality-adjustments in national accounting.

Part of the hedonic model's appeal has always been the simple relationship between hedonic prices and consumer demand: the derivative of a hedonic price function with respect to a characteristic, at a point in time, is equal to a household's marginal willingness to pay for the characteristic. This aspect of the model is highly appealing because price functions can feasibly be estimated with simple, transparent research designs, yet they also have a clear welfare interpretation. In other words, the hedonic model has something to offer both the "structural" and the "reduced form" inclinations of economists.

However, two sources of dissatisfaction with the basic hedonic model have led economists to take it in two different directions over time. First, the marginal willingness to pay potentially observed from only the hedonic price gradient is generally viewed as inadequate information for welfare evaluations of large policy shocks. Accordingly, since Rosen (1974), economists have sought to identify households' willingness-to-pay functions for amenities in a second stage based on the first-stage hedonic price function. But recovering these willingness-to-pay functions has proved a challenge. The task is complicated by the problem of unobserved demand shifters, which are systematically correlated with both levels of the characteristics obtained and marginal prices, as well as the producer characteristics to which they are matched in equilibrium (Bartik 1987, Epple 1987) (see Palmquist 2005a for a review). Proposed solutions to this problem combine, in one way or another, the economic logic of sorting along with some structure imposed on heterogeneity in tastes. In the hedonic model, Ekeland, Heckman, and Nesheim (2004) consider the case of additive hedonic models, noting that nonlinearities in the equilibrium price function justify using nonlinear functions of observed demand shifters as instruments for

the observed quantities demanded. Heckman, Matzkin, and Nesheim (2005, 2010a) and Bishop and Timmins (2015) discuss strategies for imposing functional form restrictions that allow one to map quantities of characteristics demanded into demands. Other approaches to the problem go straight to modeling the deep parameters using structural models (e.g. Bayer, Ferreira, and McMillan 2007, Sieg et al. 2004; see Kuminoff, Smith, and Timmins 2013 for a review). This literature has provided a tremendous advance on our ability to model general equilibrium counterfactuals as well as non-marginal welfare effects. However, these advantages also come at the price of imposing additional structure and complication, losing some of the simple reduced-form appeal of the hedonic model.

Second, economists increasingly have emphasized the problem of unobserved characteristics correlated with those of interest, which may bias estimates of the hedonic price function (and hence marginal willingness to pay). Accordingly, using panel data, they have applied difference-in-differences and other quasi-experimental research designs to the hedonic model (e.g. Chay and Greenstone 2005, Currie et al. 2015, Greenstone and Gallagher 2008). These approaches have the advantage of identifying a reduced form relationship between prices and characteristics under credible assumptions about the source of variation in the data. However, because they mix information from two or more hedonic equilibria, whereas the welfare measures are rooted in a single cross-sectional equilibrium relationship, these advantages come at the cost of a transparent connection to the underlying structural model and hence to the ability to make welfare inferences (Klaiber and Smith 2013, Kuminoff and Pope 2014).

In this paper, I reconsider the hedonic model in the spirit of calls from Chetty (2009) and Heckman (2010) to seek compromises that combine the clarity of reduced form econometric models with the ability of structural models to speak to welfare effects. Chetty (2009) recommends economists look for simple "sufficient statistics" that can be used to quantify non-marginal welfare measures. Heckman (2010) similarly urges us to follow "Marschak's maxim" and solve well-posed economic problems with minimal assumptions. Accordingly, I reconsider how much information we may be able to obtain from only the first-stage hedonic price function. I show that, with multiple time periods, it is possible to combine the economic logic of the hedonic model with estimation of only Rosen's first stage hedonic price function to identify non-marginal welfare effects under minimal assumptions. This is in contrast to the standard view that knowledge of the hedonic price function alone is insufficient to analyze welfare effects of large

policy shocks with general equilibrium effects.

In particular, in Section 2 of this paper, I consider the situation where it is possible to estimate hedonic price functions in a single cross-section, but where we seek more welfare information than marginal values. I assume the observed policy is the only change to the economic environment shifting implicit prices or equilibrium levels of characteristics; thus the approach taken in this section is most relevant for a narrow window surrounding a relatively sudden policy shock. In this case, the first-stage hedonic price functions can be used to derive a "sufficient statistic" for non-marginal welfare changes, in the sense of Chetty (2009). This approach provides a hybrid between simpler reduced-form and structural approaches to hedonic estimation. It simplifies the estimation problem by side-stepping the difficult endogeneity problem of Rosen's second stage or associated structural models, while using the full economics of the hedonic model to make inference about sorting and welfare effects. I also consider the minimal assumptions about household demands for this approach. In the case where panel data on households and their choices are available and where demands are constant, the sufficient statistic approach is feasible under very general conditions and heterogeneity does not need to be modeled explicitly. In the more common cases where only repeated cross-sections of hedonic prices are available or where demands are shifting, the sufficient statistic approach remains feasible under additional restrictions to heterogeneity, namely a single-crossing restriction commonly invoked in the existing literature.

In Section 3, I consider the opposite situation, where omitted variables potentially correlated with the attribute of interest motivate difference-in-differences hedonic models. While such models can overcome endogeneity problems in estimation, the hedonic literature has not reached a consensus on the interpretation of just what such models actually identify and whether they are economically meaningful. In contrast both to one recent interpretation of these models as identifying a vague "capitalization effect" that mixes information from different equilibria (Klaiber and Smith 2013, Kuminoff and Pope 2014) and to another interpretation of them as identifying a Marshallian welfare measure (Greenstone and Gallagher 2008), I show that the estimand defined by difference-in-differences hedonics has a clear interpretation as a *lower bound on the Hicksian equivalent surplus* for a non-marginal change in characteristics, even when general equilibrium effects are present. The bound is analogous to measures first suggested by Bartik (1988) and Kanemoto (1988). The bound is valid under a much wider set of condi-

tions that those of Section 2. In particular, the economic environment can change in other ways than the policy of interest and households' demands can shift between periods.

In Section 4, I demonstrate these results using simulations of hedonic equilibria. The results of the simulations are consistent with the empirical approach outlined. The sufficient statistic approach is a compromise between two Hicksian welfare measures (compensating variation [CV] and equivalent variation [EV]) and in practice it may hold even when the sufficient conditions justifying it do not. When omitted variables are of concern, estimates of welfare effects may be biased, but a lower bound measure to Hicksian equivalent surplus (ES) is still valid.

To fix ideas, I specifically discuss the example of housing markets with spatially varying amenities and I primarily will discuss connections to that literature. However, the implications of this paper are not limited to that setting and apply equally to labor markets or to other contexts with differentiated consumer products.

2. Using only the First Stage to Evaluate Non-Marginal Changes in Amenities: The Hedonic Harberger Triangle

In this section, I consider the case where the "first stage" hedonic price functions can be estimated credibly in individual cross-sections. This has been the traditional hedonic approach for decades and continues to be invoked in many models (e.g. Bishop and Timmins 2015, Ekeland, Heckman, and Nesheim 2004, Heckman, Matzkin, and Nesheim 2005, 2010a). In this setting, I reconsider how much we can learn about welfare effects for changes in amenities using only the marginal prices from the first stage, but where data are available for multiple time periods.

Since the seminal work of Rosen (1974), much of the hedonic literature has focused on the problem of estimating willingness-to-pay functions for amenities using only data from a single cross section. In this enterprise, the inherent challenge is the fact that only one point on each individual's demand function is observed, so the only variation in the data comes from the way different households sort across choice alternatives in equilibrium. The standard solution is to model heterogeneity in individual demands. Unfortunately, the unobserved components in demand (e.g. tastes) will systematically vary both with levels of amenities and their marginal prices. This correlation gives rise to the well-known endogeneity problem for Rosen's "second stage" (Bartik 1987, Bishop and Timmins 2015, Epple 1987, Heckman, Matzkin, and Nesheim 2010a).

To side-step this problem, the literature has identified two special cases in which information on non-marginal values can be obtained from only the first-stage hedonic price function. The first is the case of no heterogeneity in preferences. As Rosen (1974), noted, in this case everybody must be indifferent between all products, so the equilibrium price function traces out an indifference curve. Thus, predicted changes in prices along a constant hedonic price function represent non-marginal CV or EV for amenity improvements in partial equilibrium. A second special case is the "localized externality" of Palmquist (1992), in which exogenous changes in amenities isolated in a small part of the hedonic market would not shift the hedonic price function. Households would relocate to their initial bundle, but suppliers would be better off by the increased value of housing. In this model, even with heterogeneity, predicted changes in prices along a constant hedonic price function still represent non-marginal CV and EV.

Palmquist's (1992) model is actually a limiting case of the model discussed by Bartik (1988) and Kanemoto (1988). Bartik and Kanemoto showed that, in general, even if the hedonic price function does shift, and even if there are adjustments in housing attributes, aggregate predicted price changes using the ex ante hedonic price function represent an upper bound on welfare changes. By the same token, using their logic, it is easy to show that predicted capitalization using the ex post hedonic price function represents a lower bound.

Over the past 15 years, researchers have made important advances, moving beyond these special cases. Although they have proposed a variety of approaches to the problem, all have the common element of explicitly modeling heterogeneity in demand together with the equilibrium sorting process, whether in the continuous world of the hedonic model (e.g. Bajari and Benkard 2005, Bishop and Timmins 2015, Heckman, Matzkin, and Nesheim 2005) or the discrete-continuous world of "sorting models" (e.g. Bayer, Ferreira, and McMillan 2007, Ferreyra 2007, Kuminoff 2012, Sieg et al. 2004; see Kuminoff, Smith, and Timmins 2013 for a review). Because the economics of the models imply a particular mapping from households' preferences and incomes to the way they sort in equilibrium, researchers have suggested various ways to invert the logic and recover preferences from observed sorting. To identify the model, one approach involves imposing additional structure on the problem in the form of distributional assumptions about unobserved tastes. For example, these tastes may be assumed to have an extreme value distribution (e.g. Bayer, Ferreira, and McMillan 2007) or a log-normal distribution (e.g. Sieg et al. 2004); similarly, willingness-to-pay (WTP) functions may be assumed to have errors follow-

ing some known distribution (e.g. normal in Bishop and Timmins 2015). Alternatively, one can relax these distributional assumptions but forego point identification of the underlying parameters and be content with set identification (Kuminoff 2012).

As Bajari and Benkard (2005) and Kuminoff and Pope (2012) have pointed out, the problem becomes considerably easier when individuals are observed in multiple settings, as then individuals' willingness-to-pay functions can be fitted to two or more points. Such data are becoming increasingly available, even in the context of housing markets. For example, in the United States, researchers are beginning to make use of data available under the Home Mortgage Disclosure Act (HMDA) to match households to the houses they live in over time (e.g. Bayer, McMillan, and Rueben 2011, Bayer et al. 2012, Bishop and Timmins 2015, Depro and Timmins 2012). In principle, such data may be even easier to come by in other contexts, such as automobile or computer purchases.

To my knowledge, however, the literature has not noticed that when households are observed two or more times, their observed choices, together with knowledge of the first stage hedonic price functions, are sufficient to estimate welfare measures that are proportional to second-order approximations to a change in utility for *any* constant demand function—without explicitly modeling heterogeneity at all and without imposing any distributional assumptions on unobserved demand parameters.¹ In subsection 2.2, I first show this result.

In subsequent subsections, I weaken the information available to the analyst and/or the assumption of unchanging demand functions. In particular, in subsection 2.3, I consider the case where only panel data on housing characteristics and repeated cross sections of prices are available, but no data on how households sort. In this case, a complete ordering of households satisfying a single crossing property is required to recover the same information from repeated cross sections. In subsection 2.4, I consider the case of changing demands and find a similar result.

¹ Recently, Kuminoff and Pope (2012) suggest using exogenous shifts in the supply of amenities to derive within-market instruments for Rosen's "second stage." Although their suggestion is based on the same basic insight of this paper (that exogenous supply shocks can trace out a demand curve), in contrast I am suggesting that a similar procedure replace the second stage entirely, to identify a sufficient statistic for welfare measurement without estimating the deep structural parameters.

2.1 Primitive Notions

Throughout this paper, I consider a closed city (or region) with a constant set of households.² Let \mathcal{H} denote the set of houses with typical element h and let \mathcal{I} denote the set of households with typical element i . Equilibrium in each time period consists of a one-to-one correspondence of households to houses (all households occupy a house and all houses are occupied by a household). Households rent their houses from absentee landlords.³

Houses are differentiated by price p , the continuous amenity of interest g , and a vector of continuous housing characteristics \mathbf{x} with characteristics indexed by $r=\{1,\dots,R\}$ (lot size, dwelling size, and so forth). (The variable g may be thought of as an index of public goods, or alternatively other public goods of secondary interest may be thought to be embedded in \mathbf{x} .)

Notationally, it will sometimes be more convenient to work with a more parsimonious notation with the vector $\mathbf{z}' = [g, \mathbf{x}']$ and with the elements of \mathbf{z} indexed by j .

At any point in time t , households differ by their income y and by their current-period preferences, which can be represented by a twice differentiable quasi-concave conditional indirect utility function $v_i^t(y_i^t - p_h, g_h, \mathbf{x}_h)$, with $\partial v_i^t / \partial y_i^t > 0$ and $\partial v_i^t / \partial g \neq 0$ everywhere $\forall i$. Note that $-\partial v_i^t / \partial p = \partial v_i^t / \partial y_i^t \equiv \lambda_i^t$.

On the supply side of the market, the profit function for house h is $\pi_h = p_h - c_h(\mathbf{x}_h)$, where the cost function $c_h(\cdot)$ is twice differentiable. For convenience, I assume $c_h(\cdot)$ is constant over time, although this assumption could be relaxed.

Consider two time periods, with $t=0$ in the initial situation and $t=1$ in a later situation. Let $F^t(\cdot)$ be the distribution function of g at time t . Prices of houses are determined by the amenities and the equilibrium price function: $p_h^t = p^t(g_h^t, \mathbf{x}_h^t)$. The time superscript on the hedonic price function indicates that equilibrium hedonic prices may shift over time. In principle, these shifts may happen from changes in the distribution of g , changes in household demands, or other changes in the economic environment. In the remainder of Section 2 I will rule out the third and

² The area modeled does not literally need to be one city (or housing market). Nor need it coincide with the area affected by the policy of interest. However, as always, economists modeling demand must make judgments about the set of relevant substitutes.

³ I impose this restriction here to facilitate the exposition. Strictly speaking, this assumption is only required for the models described in Sections 2.3 and 2.4. The basic model in Section 2.2 does not require this assumption, nor does that of Section 3.

in parts of Section 2 I will rule out the second, but in general unless explicitly stated any of these sources is permissible.

In the initial situation, the household maximizes utility over a continuous choice set defined by the continuously differentiable hedonic function $p^0 = p^0(g^0, \mathbf{x}^0)$. A policy then exogenously shocks the distribution of the amenity g available in the city. Additionally, tastes and preference parameters may change between periods. As a consequence of both effects, the equilibrium price function adjusts to $p^1 = p^1(g^1, \mathbf{x}^1)$, with the set of other available characteristics \mathbf{x} possibly changing endogenously.

I make the standard hedonic assumption that households are in a static equilibrium in each time period.⁴ Maximizing utility in period t , the household satisfies the first-order condition:

$$(1) \quad \frac{\partial v_i^t}{\partial z_j} = - \frac{\partial v_i^t}{\partial p} \frac{\partial p^t}{\partial z_j} = \lambda_i^t \frac{\partial p^t}{\partial z_j}.$$

Equation (1) represents the standard tangency condition, in which the derivative of the hedonic function with respect to an amenity is equal to marginal willingness to pay for the amenity at the optimal point.

Similarly, the landlord's first-order condition for profit maximization is

$$(2) \quad \frac{\partial c_h}{\partial x_r} = \frac{\partial p^t}{\partial x_r}.$$

The endogenous amenities \mathbf{x} are supplied according to similar tangency condition, with marginal cost of supply equal to the marginal revenue.

The basic problem is to make inferences about non-marginal welfare effects from these primitive conditions.

2.2 Non-marginal values when demands are constant and a panel of households is availa-

⁴ This assumption continues to underlie the vast majority of work on hedonic markets (e.g. Bajari and Benkard 2005, Bishop and Timmins 2015, Ekeland, Heckman, and Nesheim 2004, Heckman, Matzkin, and Nesheim 2010) as well as structural sorting models of locational choice (e.g. Bayer, Ferreira, and McMillan 2007, Kuminoff 2012, and Sieg et al. 2004). However, recent work is beginning to consider dynamic optimization in the context of transaction costs, which may be substantial in applications to housing (Bayer et al. 2011, Bishop 2012, Kennan and Walker 2011). The labor literature has a longer tradition of considering such dynamic optimization (e.g. Keane and Wolpin 1997).

ble

As a starting point, consider the arguably restrictive case covered by the following three assumptions.

ASSUMPTION A1 (Panel of Households). Panel data on household choices are available, so that p and $\partial p^t / \partial z_j$ can be evaluated for each household in each time period at their choice of \mathbf{z} .

ASSUMPTION A2 (Constant Demands). Households' preferences and incomes are constant over the time period considered: $v_i^t(y_i^t - p_h, g_h, \mathbf{x}_h) = v_i(y_i - p_h, g_h, \mathbf{x}_h) \forall t$, so that the optimal vector \mathbf{x} and g are unchanging functionals of the hedonic price function.

ASSUMPTION A3 (Constant Economic Environment). The relevant policy shock to the distribution of g is the only change in the economic environment shifting slopes of the hedonic price function and the equilibrium levels of \mathbf{x} .

In this subsection, I show that, under Assumptions A1-A3, a second-order approximation of the general equilibrium welfare effects of a change in amenities can be constructed using only estimated marginal prices. The assumptions essentially guarantee that observed shifts in conditions are what is to be evaluated (A3), that these shifts trace out a demand curve (A2), and that points on the demand curve are observed (A1). Given Assumption A3, the approach taken in this section may be most relevant for sudden changes in conditions, such as discovery of a cancer cluster (Davis 2004) or the release of school report cards (Figlio and Lucas 2004). Below, Assumptions A1 and A2 will be relaxed in turn.

For consumers, the Marshallian measure of the benefits of the (non-marginal) change in the distribution of g from F^0 to F^1 is given by:

$$(3) \quad WTP_i = -dp_i + \int_{t=0}^1 \frac{\partial p^t}{\partial g} [g_i^*(t), \mathbf{x}_i^*(t)] dt,$$

where $\mathbf{x}_i^*(t)$ and $g_i^*(t)$ are the household's optimal levels of the respective amenities at notional time t given the prevailing price function; where $g_i^*(t)$, $\mathbf{x}_i^*(t)$, and the price function continuously adjust between $t=0$ and $t=1$; and where dp_i is the price change experienced by the household (Scotchmer 1986, Bartik 1988). Equation (3) is impossible to observe literally, but it is instructive. It reduces the problem of measuring non-marginal willingness to pay to an index number problem, that is, to an average of marginal willingness to pay along the path between $[p^0, g^0, \mathbf{x}^0]$ and $[p^1, g^1, \mathbf{x}^1]$.

Although Equation (3) is based on a Marshallian construct, Arnold Harberger famously suggested that a linear approximation to (3) could be interpreted as a valid approximation to an exact welfare measure.⁵ More recently, Chetty (2009) has suggested that Harberger's approach can be thought of as setting a paradigm for sufficient-statistic welfare measurement. Following Harberger (1971), consider a second order approximation to a change in utility for an individual in the hedonic model given Assumptions A2 and A3:

$$(4) \quad dv_i \approx \frac{\partial v_i}{\partial p} dp_i + \sum_j \frac{\partial v_i}{\partial z_j} dz_{j,i} + \frac{1}{2} \frac{\partial^2 v_i}{\partial p^2} dp_i^2 + \sum_j \frac{\partial^2 v_i}{\partial p \partial z_j} dp_i dz_{j,i} + \frac{1}{2} \sum_{j'} \sum_j \frac{\partial^2 v_i}{\partial z_j \partial z_{j'}} dz_{j,i} dz_{j',i},$$

where $dp_i = p^1(g_i^1, \mathbf{x}_i^1) - p^0(g_i^0, \mathbf{x}_i^0)$ is the household's change in expenditure and $dz_{j,i}$ is the change in amenity j experienced by household i after all adjustments. These changes stem from a number of sources. At the household's initial optimal location, g may change directly from the policy and the price of the home may capitalize this change. Additionally, p changes as the hedonic function shifts. Finally, p , g , and \mathbf{x} may all change from any readjustments by the household as it re-optimizes, and \mathbf{x} also may change from any supply-side investments as landlords re-optimize. Whatever the source of the changes, welfare effects are evaluated taking all of them into account.

Equation (4) leads to the following lemma.

LEMMA 1. Given Assumptions A1-A3, a second order approximation to the change in welfare for each consumer, dw_i , from an exogenous change in the distribution of g , can be constructed from observed prices and estimated marginal prices as follows:

$$(5) \quad dw_i \equiv \frac{dv_i}{\frac{1}{2}(\lambda_i^0 + \lambda_i^1)} \approx -dp_i + \sum_j \frac{1}{2} \left(\frac{\partial p^0}{\partial z_j} \Big|_{z_i^0} + \frac{\partial p^1}{\partial z_j} \Big|_{z_i^1} \right) dz_{j,i}.$$

Proof: See the appendix.

This expression is proportional to the utility change dv , which is converted to the measuring rod of money using the average marginal utility of income, averaged between the starting

⁵ See Banzhaf (2010) for a discussion of this approach to welfare measurement in a historical context.

point and ending point. Lemma 1 states that this change in welfare for a consumer is given by the change in rents dp , plus the change in housing attributes and public goods experienced by the household after all adjustments, multiplied by the *average* marginal willingness to pay, again averaged between the starting point and ending point. The expression might be thought of as a "hedonic Harberger triangle" (or trapezoid). It is a compromise between two Hicksian measures, CV and EV.

On the supply side of the market, landlords are directly better off by the change in rents dp . This change in rents stems from shifts in the price function and from exogenous changes in g , but also potentially from adjustments to \mathbf{x} that are costly to supply. Consequently, the cost of producing the change in \mathbf{x} must be netted out of the change in profits. The change in profits from any change in g , the price function $p(\cdot)$, or endogenous adjustments to \mathbf{x} is $d\pi = dp - dc$. We can in turn take a second-order approximation to dc as follows.

$$(6) \quad dc_h \approx \sum_r \frac{\partial c_h}{\partial x_r} dx_{r,h} + \frac{1}{2} \sum_{r'} \sum_r \frac{\partial^2 c_h}{\partial x_r \partial x_{r'}} dx_{r,h} dx_{r',h}$$

This fact along with the first-order conditions leads to the following lemma.

LEMMA 2. A second order approximation to the change in profits for each landlord, $d\pi_h$, can be constructed from observed prices and estimated marginal prices as follows:

$$(7) \quad d\pi_h \approx dp_h - \sum_r \frac{1}{2} \left(\frac{\partial p^0}{\partial x_r} \Big|_{z_h^0} + \frac{\partial p^1}{\partial x_r} \Big|_{z_h^1} \right) dx_{r,h}.$$

Proof: See the appendix.

LEMMA 2 says that the change in profits is just the change in price, net of an adjustment accounting for changes in costs due to endogenous changes in \mathbf{x} , which can be approximated from marginal prices.

Let the change in aggregate welfare W be given by aggregating over the changes in consumer surplus and profits:

$$(8) \quad dW \equiv \int_j dw_i di + \int_{\mathcal{H}} d\pi_h dh.$$

By lemmas 1 and 2, we can integrate over Expressions (5) and (7) and substitute them into the respective terms in Equation (8). Additionally, we can combine these into one integral, but doing so requires some additional notation because Expression (5) is evaluated at the choices made by a single household i (regardless of location) whereas Expression (7) is evaluated at a particular house h (regardless of who lives there). Let $\mathbf{z}_{i(\tau)}^t$ represent the characteristics, in time t , of a house actually occupied by household i in time τ . Similarly, let $\frac{\partial p^t}{\partial \mathbf{z}_j} |_{\mathbf{z}_{i(\tau)}^t}$ be the partial derivative of the period t hedonic price function with respect to attribute j , evaluated at the period t attributes of the house actually occupied by household i in time τ . By definition, $\mathbf{z}_{i(t)}^t = \mathbf{z}_i^t$. However, this more general notation allows us to keep track of a household's former house or future house, even when it is not currently living there.

Equation (8) together with this notation lead to the following proposition.

PROPOSITION 1. Given Assumptions A1-A3, a second-order approximation to the change in aggregate welfare from an exogenous change in the distribution of g , when prices, households, and landlords adjust to the change endogenously, is

$$(9) \quad dW \approx \int_J \frac{1}{2} \left[\left(\frac{\partial p^0}{\partial g} |_{\mathbf{z}_{i(0)}^0} + \frac{\partial p^1}{\partial g} |_{\mathbf{z}_{i(1)}^1} \right) dg_i + \sum_r \left(\frac{\partial p^0}{\partial x_r} |_{\mathbf{z}_{i(0)}^0} + \frac{\partial p^1}{\partial x_r} |_{\mathbf{z}_{i(1)}^1} \right) (x_{r,i(1)}^1 - x_{r,i(0)}^1) + \sum_r \left(\frac{\partial p^1}{\partial x_r} |_{\mathbf{z}_{i(1)}^1} - \frac{\partial p^1}{\partial x_r} |_{\mathbf{z}_{i(0)}^1} \right) (x_{r,i(0)}^1 - x_{r,i(0)}^0) \right] di.$$

Proof: The proposition follows from Lemmas 1 and 2, simply integrating (5) over J and (7) over \mathcal{H} , adding the two together, fixing our indices so $h = i(0)$ (which we can do given the bijective mapping between them), and re-arranging terms.⁶

Despite the intimidating subscript notation, Expression (9) is actually quite simple. It consists of three terms. The first is the change in g experienced by a household (across houses if it moves) evaluated at the partial derivative of the hedonic price function, averaged over those

⁶ The first pair of terms in parentheses in Expression (9) comes from Expression (5). Adding the other terms of Expressions (5) and (7), we have $\frac{1}{2} \sum_r \left(\frac{\partial p^0}{\partial x_r} |_{\mathbf{z}_{i(0)}^0} + \frac{\partial p^1}{\partial x_r} |_{\mathbf{z}_{i(1)}^1} \right) dx_{r,i} - \frac{1}{2} \sum_r \left(\frac{\partial p^0}{\partial x_r} |_{\mathbf{z}_{i(0)}^0} + \frac{\partial p^1}{\partial x_r} |_{\mathbf{z}_{i(0)}^1} \right) dx_{r,i(0)}$. We can write $dx_{r,i} = (x_{r,i(1)}^1 - x_{r,i(0)}^1) + (x_{r,i(0)}^1 - x_{r,i(0)}^0) = (x_{r,i(1)}^1 - x_{r,i(0)}^1) + dx_{r,i(0)}$. The rest follows by regrouping terms.

two points. The second term is the change in \mathbf{x} experienced by the household, accounting for only the fact that it may move (i.e. netting out any supply adjustments in \mathbf{x} at a fixed location), but with these changes again evaluated by the respective average partial derivatives experienced by the household (across houses). The third and final term is the supply response at the household's initial house, multiplied by the difference in the ex post partial derivatives between the final location and the initial location. The changes in prices dp which appear in Expressions (5) and (7) cancel out as transfers between households and landlords.

Proposition 1 is the key result of this sub-section. It says that, under our assumption that we observe a panel of housing attributes and our initial assumption that we can track households across locations, a second-order approximation to welfare can be estimated with only the derivatives of the first-stage hedonic price function at the two points in time. The approach essentially relies on the fact that the policy shock induces a new equilibrium, which allows demands to be approximated by tracing out the optimized points.

Expression (9) can be simplified if we are willing to assume that the supply side of the market is highly competitive and that profits (and rents to fixed factors) from adjustments to \mathbf{x} are approximately zero. Denote this as Assumption A4:

ASSUMPTION A4 (Zero profits). The change in profits due to adjustments in \mathbf{x} are approximately zero: $\int_h [p^1(g_h^1, \mathbf{x}_h^1) - p^1(g_h^1, \mathbf{x}_h^0)] dh \approx \int_h [c(g_h^1, \mathbf{x}_h^1) - c(g_h^1, \mathbf{x}_h^0)] dh$.

Thus, the change in profits consists only of the change in rents to land and to the sunk capital \mathbf{x}^0 :

$$\int_h d\pi dh \approx \int_h (p^1(g^1, \mathbf{x}^0) - p^0) dh = \int_h [dp - (p^1 - p^1(g^1, \mathbf{x}^0))] dh.$$

In this case, we would have the following variant of Proposition 1:

PROPOSITION 2. Given A1-A4, a second-order approximation to the change in aggregate welfare from an exogenous change in the distribution of g , when prices, households, and landlords adjust to the change endogenously, is:

$$(10) \quad dW \approx \int_j \sum_i \frac{1}{2} \left(\frac{\partial p^0}{\partial z_j} \Big|_{z_i^0} + \frac{\partial p^1}{\partial z_j} \Big|_{z_i^1} \right) dz_{j,i} di - \int_{\mathcal{H}} (p_h^1 - p_h^1(\mathbf{x}^0, g^1)) dh.$$

Proposition 2 states that we can measure benefits by tracking the average change in amenities experienced by each household, weighted by the average of the derivatives of the two hedonic price functions evaluated at the households' choice, netting out any aggregate price

changes that reflect real costs of adjustments in \mathbf{x} . For simplicity, I will use Expression (10) as the measure of benefits for the remainder of Section 2. However, this choice plays no substantive part of the results later derived, and parallels could be derived based on Expression (9) as well.

Note finally that in the simple case where there are no supply adjustments, then $x_{i(0),r}^1 = x_{i(0),r}^0$ and the entire expression collapses to

$$(11) \quad dW \approx \int_j \sum_j \frac{1}{2} \left(\frac{\partial p^0}{\partial z_j} \Big|_{z_i^0} + \frac{\partial p^1}{\partial z_j} \Big|_{z_i^1} \right) dz_{j,i} di.$$

Expression (11) is just the individual household measure from Expression (5) summed over households, but with the dp terms cancelling as transfers between residents and landlords.

Figure 1 illustrates measure (11), or equivalently the first terms in Expressions (9) and (10). The figure shows the derivatives with respect to g of both the before-policy and after-policy hedonic functions; these derivatives are positive and continuous but unrestricted as to slope or curvature. The line $\text{bid}(g)$ represents the linearized approximation to the Marshallian bid function. The points g^0 and g^1 represent the levels of the public good selected by the consumer in each scenario. Although it is a Marshallian measure, the area under the linearized bid function represents a second-order approximation to the welfare change associated with this change in g . The figure illustrates this measure only in the dimension of g , but note Expressions (10) and (11) require summing over all attributes j . Even if there are no adjustments to \mathbf{x} in the housing stock as a result of the policy, i.e. even if $\int_j d\mathbf{x}_i di = \mathbf{0}$ and only g changes in aggregate, the welfare measure for this change in g still requires taking these terms into account, weighted by the average marginal WTP of the household, as shown in Expressions (10) and (11). The measures are no different if people are owner-occupiers. In that case, the wealth effects still cancel and the $d\mathbf{z}$ terms incorporate the wealth effects on demand for attributes, as would be appropriate. Regardless, a valid welfare measure can be obtained simply by adding up experiences changes in characteristics, weighted by the average marginal values.

2.3 When only a panel of *houses* is available

The previous subsection involves Assumption A1, that panel data on households' choices are

available. This subsection considers what we can learn if we relax this assumption.

To see the role of Assumption A1, note that Propositions 1 and 2 introduce expressions involving $\sum_j \frac{1}{2} \left(\frac{\partial p^0}{\partial z_j} \Big|_{z_i^0} + \frac{\partial p^1}{\partial z_j} \Big|_{z_i^1} \right) dz_{j,i}$ for each household i . This expression involves observing (i) in which house each household lives in both time periods; (ii) a panel of house attributes (g, \mathbf{x}) or, if we are willing to assume $d\mathbf{x}=0$, a panel of g together with data on \mathbf{x} at one time period; and (iii) either a panel of house prices (as in a repeat sales model) or repeated cross sections of housing prices sufficient for predicting p_h at each location. Of these, the first is the most difficult to observe: in many cases analysts may observe housing characteristics and a sample of prices, but have no information whatsoever on who is occupying those houses. Unfortunately, in general such data are not sufficient to use the strategy outlined in the previous sub-section. Expression (10) requires information on the *household's* change in amenities, dz_i . In general, information on the change in amenities at a fixed house h , dz_h , are not sufficient. The problem is depicted in Figure 2. The figure shows two points respectively chosen by two households in the first period, two in the second, and two possible pairs of demands for the two households. However, which household sorts into which house is unobserved. Consequently, the two households' demand curves may be the two solid lines or alternatively the two dashed lines. But the sums of the areas under those respective pairs of demand curves are not the same.

However, information on only the changes in g and other characteristics of the houses, which is available in a great many hedonic applications (perhaps the majority), is sufficient under at least two special cases. The first is trivial: if households do not relocate in equilibrium then one can substitute the change in amenities at a fixed house h , dz_h , for the change in amenities consumed by a household, dz_i , in Expression (10). In this case, aggregating over houses amounts to the same thing as aggregating over households. However, for the large changes envisioned here, this patterns is unlikely to hold.

Nevertheless, data on dz_h can still be used under a second, more interesting, case: a "single-crossing" restriction on any two households' indifference curves. This amounts to a restriction on preferences such that households' Marshallian bid functions for some observed amenity do not cross. That is, households can be ordered by their marginal willingness to pay for the amenity, and the ordering will be the same evaluated at any level of the amenity and under any equilibrium price function. Because households always sort in the same order, if we

have single crossing on g then in Figure 2 we can infer that the sorting is that of the dashed lines and rule out the solid lines. Essentially, the logic of single crossing provides a way to impute households' pattern of sorting, even when their actual locations are not observed. We state this property as Assumption A5:

ASSUMPTION A5 (Single Crossing). Let $V_{ji}(y_i, p(\cdot), z_j)$ be the indirect function conditional only on attribute z_j given the price function $p(\cdot)$, with the other attributes optimally chosen subject to z_j , y , and $p(\cdot)$ to determine the utility level. Let \mathcal{J} be a simply ordered set and let the distribution of demands be such that, for some amenity z_j , $v_i^t / \partial z_j \neq 0 \forall i$ and $\frac{\partial v_{j,i}(y_i, p(\cdot), z_j) / \partial z_j}{\partial v_{j,i}(y_i, p(\cdot), z_j) / \partial y}$ is everywhere non-decreasing in i .

Assumption A5 requires single crossing in only one dimension of the characteristic space.⁷ The public good of interest g may be a natural choice for that attribute, but that choice is not necessary. z_j could be any characteristic or any scalar-valued index of characteristics. Even though we are modeling multidimensional characteristics, induced assortative matching along any one dimension is enough to impute households' choices. The single-crossing assumption, sometimes called the Spence-Mirrlees condition when written this way, guarantees that the level of z_j chosen by households is always increasing in i .⁸ See, e.g., Milgrom and Shannon (1994) and Athey, Milgrom, and Roberts (1998 Ch. 3) for proofs.⁹ To see this intuitively, note that we could write A5 alternatively as follows. For any two households i, i' if $\frac{\partial v_{j,i}(y_i, \hat{p}(\cdot), \hat{z}_j) / \partial \hat{z}_j}{\partial v_{j,i}(y_i, \hat{p}(\cdot), \hat{z}_j) / \partial y} \geq$

⁷ As Chiappori, McCann, and Nesheim (2010) discuss, extensions of the single-crossing property to the multi-attribute case lose the interpretation of inducing assortative matching. Although we are working with a multi-attribute model, Assumption 5 involves single-crossing of the WTP functions in the dimension of only one attribute.

⁸ As discussed by Athey, Milgrom, and Roberts (1998), Edlin and Shannon (1998), and Milgrom and Shannon (1994), a more general version of single crossing using monotone comparative statistics is sufficient to guarantee sorting by i . However, in the hedonic context we are already assuming the regularity conditions, associated with the hedonic tangency conditions, which as they show imply the Spence-Mirrlees condition specified here. (In particular, the assumptions that $v(\cdot)$ is continuously differentiable, $\partial v / \partial y > 0$, and $\partial v / \partial z_j \neq 0$.) If we were to relax the tangency conditions and focus on welfare bounds associated with inequality conditions instead of equalities, the required restriction would be their single crossing condition (see e.g. Edlin and Shannon 1998, condition 1).

⁹ See, e.g., Theorem 3.2 in Athey, Milgrom, and Roberts (1998). Their proof applies directly if we assume the utility-maximizing choice for each household is a singleton. Moreover, it is easy to see that, with single crossing, a household can only be indifferent between two points on an equilibrium hedonic price function if at least one other household is too (otherwise, another household would bid up the price of the house it prefers and the first household would no longer be indifferent). But such ties are precisely the case where mistaking the sorting is immaterial.

$\frac{\partial v_{j,i'}(y_{i'}, \hat{p}(\cdot), \hat{z}_j) / \partial \hat{z}_j}{\partial v_{j,i'}(y_{i'}, \hat{p}(\cdot), \hat{z}_j) / \partial y}$ for some $\hat{z}_j, \hat{p}(\cdot)$, then $\frac{\partial v_{j,i}(y_i, p(\cdot), z_j) / \partial z_j}{\partial v_{j,i}(y_i, p(\cdot), z_j) / \partial y} \geq \frac{\partial v_{j,i'}(y_{i'}, p(\cdot), z_j) / \partial z_j}{\partial v_{j,i'}(y_{i'}, p(\cdot), z_j) / \partial y}$ for all $z_j, p(\cdot)$.

Thus, if a household selects more z_j than another household in the baseline scenario, it will do so in the ex post scenario as well. Intuitively, given that the Marshallian demand functions do not cross, this obviously must be so if the implicit price of z_j is increasing in z_j (i.e. if the hedonic price function is convex in z_j). However, even if the implicit price is decreasing in z_j over portions of the range, the second-order condition for utility maximization requires that it cut the demand curves from below. In other words, the economics of the model require that the slope of the price function be greater than the slope of the demand curve in the neighborhood of the optimal choice. Thus, households will always "sort" across z_j in the same order, even if they are changing consumption of other attributes or the numeraire.

All this suggests a simple approach for identifying Expression (10) with panel data on houses. Let \mathcal{H} and \mathcal{I} now be finite countable sets indexed by $h = \{1, \dots, H\}$ and $i = \{1, \dots, I\}$ with $H = I$. These can be viewed as finite samples of data drawn from the underlying distribution. Let $F_j^t(\cdot)$ be the distribution function of some continuously distributed amenity z_j in period t . Given that z_j is continuously distributed, for each observed percentile $\theta \in \{1/H, 2/H, \dots, 1\}$ of the distribution of z_j , there will be a unique vector $\mathbf{z}^t(\theta)$ in period t . Let $z_k^t(\theta)$ be the value of the k^{th} attribute of this vector. Note for $k=j$, $z_k^t(\theta) = (F_j^t)^{-1}(\theta)$. Then we can now state the following proposition.

PROPOSITION 3. Under Assumptions A2-A5, aggregate welfare for a change in the distribution of g , when prices, households, and landlords adjust to the change endogenously, can be computed from observed prices and estimated derivatives as follows:

$$(12) \quad dW \approx \sum_{\theta} \sum_k \frac{1}{2} \left(\frac{\partial p^0}{\partial z_k} | \mathbf{z}^0(\theta) + \frac{\partial p^1}{\partial z_k} | \mathbf{z}^1(\theta) \right) [z_k^1(\theta) - z_k^0(\theta)] - \sum_h (p_h^1 - p_h^1(\mathbf{x}^0, \mathbf{g}^1))$$

Proof: Consider the initial equilibrium described by the hedonic price function $p^0(g, \mathbf{x})$. For any $\theta \in [1/H, 2/H, \dots, 1]$, consider a household which chooses z_j in the initial equilibrium such that $F_j^0(z_j) = \theta$. By Assumption A5 (single crossing), when facing the new equilibrium price function $p^1(g, \mathbf{x})$, the household would choose z_j such that $F_j^1(z_j) = \theta$. The remainder follows by Proposition 2, simply evaluating each household at the \mathbf{z} corresponding to a constant percentile of the distributions of z_j .

Proposition 3 states that one can rank the houses by z_j in both scenarios, find the change in each attribute at a constant percentile of the z_j distribution, and evaluate the derivatives of the hedonic price function at the \mathbf{z} falling at the same percentile of the z_j distribution. The final term is the same price adjustments as in Expression (10). In practice, note that as long as one knows g and the other attributes at all locations and time periods, this estimate can be implemented with only a repeated cross section of housing prices (i.e. without a full panel): all that is required is that p^1 and $\frac{\partial p^t}{\partial z_j}$ can be predicted for each house from the hedonic pricing model.

This single-crossing condition is routinely imposed in the literature on non-linear pricing (e.g. Athey, Milgrom, and Roberts 1998, Wilson 1993), including models of monopoly screening as well as locational sorting. Although imposing this property is undoubtedly a restriction relative to the more general treatment of heterogeneity in Section 2.2, it is actually less restrictive in this respect than many structural models, which employ the same single crossing property plus additional functional form restrictions or parametric restrictions on the distribution of unobservable demand shifters. For example, consider the common class of models which allow households i to differ in two dimensions, income y and a parameter α reflecting tastes for g . Many hedonic and sorting applications, including Bajari and Benkard (2005), Bishop and Timmins (2015), and Heckman, Matzkin, and Nesheim (2010b) impose the additional restriction that willingness to pay is strictly increasing in α and that preferences are quasi-linear. These models implicitly impose Assumption A5: households are totally ordered by α , with increasing α implying increasing g . The same is true, after taking expectations over the additive errors, of many logit models such as Bayer, Ferreira, and McMillan (2007).

This is not to say that all papers impose these conditions. Other models in this class impose only a partial ordering on i rather than a total ordering. For example, Epple, Peress, and Sieg (2010), Kuminoff (2012), and Sieg et al. (2004) order households by α conditional on y , and vice versa. However, pairs of households differing in α and y need not be ordered: one household may choose higher g than the other household in one equilibrium but not necessarily in another equilibrium. In this sense, Assumption A5 is stronger than the related single crossing assumptions imposed in those papers. However, in other respects the approach of this section still imposes weaker assumptions about heterogeneity. Epple, Peress, and Sieg (2010), Kuminoff (2012), and Sieg et al. (2004) essentially compensate for their weaker single crossing assump-

tions by imposing additional functional form restrictions on $v(\cdot)$ and (in the case of Sieg et al.) parametric assumptions on the joint distribution of (α, y) .

In practice, Expression (12) may be a reasonable approximation to Expression (10) even when households are not totally ordered, in violation of Assumption A5. In the simulations reported in Section 4, the matching-by-percentiles approach of Expression (12) gives results quite close to the matching-by-households approach of Expression (10), even when households are not totally ordered. The reason appears to be that even if households do not sort on g in exactly the same order in different scenarios, their rank orderings are still highly correlated. Consequently, violations of Assumption A5 are only local, over a range in which the hedonic price function is approximately linear (so marginal prices are approximately constant) and the households are similar enough that imputing one's marginal value to another creates only small errors. In this sense, estimates based on Expression (12) can be thought of as *approximations* to Expression (10), even when Assumption A5 does not strictly hold.

Moreover, there may be compromises between Propositions 2 and 3. The model of Section 2.2 and Proposition 2 considered the case where individual households were observed in multiple time periods. So far, this subsection and Proposition 3 have considered the polar opposite case where no information was available on the identity of which households sort into which houses. In between these two cases are a variety of intermediate ones where partial information is available on the households locating at a house. Not surprisingly, such partial information would allow us to partially relax Assumption A5.

For example, suppose we can track where *types* of households live but not the individual household. Perhaps we can observe the race of a household occupying a given house, or its income, or some other characteristic or combination of characteristics. Then we would only require that the single crossing property hold within observable type.

In this context, the required single-crossing condition may be modified as follows.

ASSUMPTION A5' (Single crossing within type). Let \mathcal{T} be a set of observable types which partitions the set of households \mathcal{H} . Let $\tau \subset \mathcal{T}$ be the set of households of a specific type with measure μ_τ , $\sum_\tau \mu_\tau = 1$, and let each τ be a simply ordered subset of the partially ordered set \mathcal{J} . For each τ , let the distribution of demands be such that, for some amenity z_j , $\frac{\partial v_{j,i}(y_i, p(\cdot), z_j) / \partial z_j}{\partial v_{j,i}(y_i, p(\cdot), z_j) / \partial y}$ is everywhere nondecreasing in $i \in \tau$.

This property guarantees that, after conditioning on the observed type, households sort along z_j in the same order. Thus, if a household selects more g than another household of the same type in the baseline scenario, it will do so in the ex post scenario as well. This approach allows households to be partially ordered but totally ordered within type.

In this case, Expression (12) can be modified as follows. Let I_τ be the number of observed households of type τ and let $\theta_\tau = \{1/I_\tau, 2/I_\tau, \dots, 1\}$ be the observed percentiles of the distribution of z_j among type τ . Then by the same argument as given in Proposition 3, aggregate welfare is

$$(13) \quad dW \approx \sum_{\tau \in \mathcal{J}} \frac{\mu_\tau}{I_\tau} \sum_{\theta_\tau} \sum_k \frac{1}{2} \left(\frac{\partial p^0}{\partial z_k} | \mathbf{z}^0(\theta_\tau) + \frac{\partial p^1}{\partial z_k} | \mathbf{z}^1(\theta_\tau) \right) [z_k^1(\theta_\tau) - z_k^0(\theta_\tau)] \\ - \sum_H (p_h^1 - p_h^1(\mathbf{x}^0, \mathbf{g}^1))$$

That is, for each type, one can take the set of houses occupied by that type and rank them by z_j in both scenarios, find the change in z_k at a constant percentile of the distribution, and evaluate the derivatives of the hedonic price function at the same percentile of the distribution. Then, one can take the weighted sum over types.

2.4 When Household Demands Shift

The models of Sections 2.2 and 2.3 both have relied on Assumption A2, that households' preferences and incomes are constant over the time period considered, so that the optimal vector \mathbf{z} is an unchanging function of the hedonic price function. As depicted in Figures 1 and 2, this assumption allows us to identify two points on a single demand vector for (g, \mathbf{x}) as a function of the hedonic prices.

Assumption A2 may be most fitting for hedonic applications to consumer goods such as computers (Bajari and Benkard 2005) or groceries (Griffith and Nesheim 2013). It may be less fitting for applications to housing if income or taste changes are important in that context or—to put it differently—when there are important changes in parameters affecting current-period utility, such as family status. In such cases, even when we observe points for two equilibria for a single household (Section 2.2), these may be two points on two different demand curves. Likewise, even if household preferences exhibit the single crossing property (Section 2.3), if incomes

or preferences are shifting across equilibria then there is no reason to believe households are still sorting in the same order.

In this subsection, I consider relaxing Assumption A2 as well as Assumption A1. Changes in the willingness to pay for g of course will have implications for equilibrium hedonic prices. But more importantly, they raise fundamental questions about what perspective to take when evaluating changes in g . I take the approach of Fisher and Shell (1972) and Pollak (1989), who suggest that welfare comparisons can be made in such situations from the perspective of one or the other preference relationships.¹⁰ This basic approach also is implicit in simulation exercises such as those in Bayer et al. (2011), Sieg et al. (2004), or elsewhere, where general equilibrium welfare effects are estimated for policy shocks under an assumption of constant preferences, even though preferences may well change over the time periods envisioned. Essentially, these simulations evaluate the policy shocks from the perspective of ex ante preferences.

From this perspective, we can again use the single crossing property along with a weaker version of Assumption A2:

ASSUMPTION A2'. The distribution of pairs $(v^t(), y^t)$ is constant over time. Consequently, the distribution of demand functionals $\mathbf{z}(p^t())$ is constant.

Assumption A2' is implied by A2 but the reverse is not true. It says that although individual households' demand functions may change over time, the distribution of demand types remains constant.

With this weaker assumption we can now state Proposition 4.

PROPOSITION 4. Under Assumptions A2' and A3-A5, aggregate welfare for a change in the distribution of g , evaluated from the perspective of constant preferences for each household, can be computed from observed prices and estimated derivatives as given by Equation (12). Moreover, the aggregate evaluation is the same whether using ex ante or ex post preferences (though the distributional effects may be different).

Proof: See the appendix.

Proposition 4 says that we can still use the same measure for aggregate welfare as given in Proposition 3 under a weaker version of Assumption A2, in which individual household demands may shift due to changes in preferences or incomes, as long as the distribution of de-

¹⁰ See Stapleford (2011) for discussion in historical context.

mands does not change. However, the welfare estimates are now from a single period perspective, or alternatively counterfactual estimates "as if" preferences had remained constant. If Assumption A2' is still considered too strong an assumption, alternative approaches may be possible introducing semiparametric controls for observed factors that change over time, such as family status (marital status, presence of children, etc.) or income. Extending the model in this direction would bring it close to an alternative model recently proposed by Bishop and Timmins (2015), which implicitly imposes single crossing plus constant distributions of demands, but imposes additional structure to condition on income and an annual additive shocks to WTP. In general, the strategy can be thought of as an across-time variant of the argument for using multiple markets under assumptions of constant distributions of tastes (Epple 1987, Heckman, Matzkin, and Nesheim 2010a).

This section has used the full economic logic of the hedonic model to derive approximations for non-marginal WTP that are a second-order approximation to any set of preference relations satisfying the respective restrictions. Proposition 2 requires constant demands, but otherwise no such restrictions on preferences beyond differentiability and monotonicity in y , but it requires observing individual households. Proposition 3 requires a single crossing restriction for any two households, but no information on the households or their type. Proposition 4 relaxes the assumption of constant demands in Proposition 2 and replaces it with an assumption of constant demand distributions and single crossing. With these assumptions, the expressions in the propositions represent "sufficient statistics" for policy evaluation using the first stage of the hedonic model, in a way that balances economic structure with generality, as urged by Chetty (2009) and Heckman (2010).

3. What Can We Learn from Difference-in-Differences Capitalization Effects?

3.1 Introduction to difference-in-differences capitalization

The models in the previous section require knowledge of the first-stage hedonic price function in two or more time periods. That is, they require estimating multiple cross sectional hedonic price functions. While working with cross-sections has been the standard hedonic project for decades, more recently economists have drawn attention to the problem of unobserved amenities that may be correlated with g . To overcome this endogeneity problem, they have applied difference-in-

differences approaches to the hedonic econometric model or introduced lagged dependent variables into the hedonic equation; they also have introduced quasi-experimental methods in which the changes to the hedonic amenities of interest are plausibly exogenous. Examples include Chay and Greenstone (2005), Currie et al. (2015), Davis (2004, 2011), Figlio and Lucas (2004), Greenstone and Gallagher (2008), Lavaine (2014), and Linden and Rockoff (2008). Parmeter and Pope (2013) provide an introduction to and review of this literature. In simulations, Kuminoff, Parmeter, and Pope (2010) illustrate the importance of using generalized difference-in-differences hedonic techniques to control for time invariant unobservables.

From a standpoint of statistical estimation, this work is clearly an important improvement, allowing for identification of the hedonic price function under much weaker assumptions about unobservable neighborhood and/or housing characteristics. However, from an economic perspective, it seemingly has come at the cost of a clear interpretation of the estimand: what economic question it answers is not always clear, or at least has not been perceived clearly in the literature. The ambiguity arises because the hedonic equilibrium is fundamentally cross sectional: households face a tradeoff among houses across space at a point in time, not across time. At a point in time, the derivative of a price function with respect to the public good, $\partial p^t(g^t, \mathbf{x}^t, \varepsilon^t)/\partial g$, is the marginal willingness to pay (where to emphasize the estimation issues we now make effects unobservable to the econometrician, ε^t , explicit in the model). In contrast, the dependent variable in a difference-in-differences hedonic is

$$(14) \quad dp = p^1(g^1, \mathbf{x}^1, \varepsilon^1) - p^0(g^0, \mathbf{x}^0, \varepsilon^0),$$

which mixes information from two equilibria. Recently, the literature has begun to refer to such differences as "capitalization" (Chay and Greenstone 2005, Klaiber and Smith 2013, Kuminoff and Pope 2014, Muehlenbachs, Spiller, and Timmins 2012, Parmeter and Pope 2013), in contrast to the slope of a single equilibrium hedonic price function.

The link between such capitalization and the underlying economic model is not immediately clear. For changes to the set of amenities for a small subset of houses, the equilibrium hedonic price function can be taken as constant over a short time period, so that difference-in-differences models can be interpreted within a single equilibrium (Palmquist 1992). But in the more general case, a large change in the supply of an amenity will shift the price of the amenity

as well as of substitutes, thus shifting the hedonic price function.¹¹ Too, other changes in the economic environment taking place over the longer time periods used in many studies, such as ten years, would also shift the price function. When the hedonic price function shifts for either reason, capitalization studies compare prices at two different equilibria in potentially confusing ways. The confusion is only compounded by ambiguity about the meaning of language, borrowed from the program evaluation literature, such as *the* "causal effect" or "capitalization effect" of a change in amenities, and about just what the implicit counterfactual might be when defining these "effects."

Consider the wide range of claims made in the literature. On one hand, Greenstone and Gallagher (2008) argue that capitalization effects represent the Marshallian consumer surplus for a change in amenities when housing is inelastically supplied (see their Figure 1 and associated discussion). However, their argument is based on two strong assumptions. First, they assume that the policy induces a *parallel shift* in the Marshallian demand for land and/or housing in the improved area. Clearly, this is a strong assumption which will not hold in general. Second, they implicitly assume that the change in Marshallian consumer surplus for housing is equal to the consumer surplus for the change in the underlying amenity. In fact, this equality does not hold except under the special case of a restriction to income effects known as the Willig condition (Palmquist 2005b and Smith and Banzhaf 2004).

On the other hand, Klaiber and Smith (2013), and Kuminoff and Pope (2014) have argued that, because it combines two equilibria, the capitalization effect answers an ill-defined economic question, selected as the estimand of interest more because there is an unbiased estimator to recover the parameter than because it is an economic parameter of interest (see also Parmeter and Pope 2013). Their argument essentially involves two points in the presence of changes in the hedonic price function. First, the total causal effect of an improvement, defined by Equation (14) when there are no other changes in the economic environment, is not the same as willingness to pay, or indeed related to it in any clear way. Instead, it conflates willingness to pay (defined within the context of an equilibrium price function) together with changes in the equilibrium price function. Second, estimation of $p^t(g^t, \mathbf{x}^t, \varepsilon^t)$ will be biased if the general equi-

¹¹ Heckman, Lochner, and Taber (1998) discuss the same issue in the context of labor market outcomes and the supply of college education.

librium effects on the price function are ignored. Both points are correct. Nevertheless, the conclusion that capitalization studies provide no meaningful welfare measure is incorrect.

In this section, I show that difference-in-differences capitalization studies identify a measure closely related to those discussed many years ago by Bartik (1988) and Kannemoto (1988), a measure which is a lower bound on ES for an improvement in g , even in the presence of shifts in the price function and adjustments to the supply of \mathbf{x} . These results are more general than those of Section 2. In particular, they do not require Assumptions A1-A3 or A5: Demands and other aspects of the economic environment may change between periods, there are no restrictions on heterogeneity in demands, and panel data on household choices are not required.

3.2 Defining capitalization effects

If the equilibrium hedonic price function for a particular housing market changes endogenously because of a shock to amenities, then the price of a house will change even if its amenities have not. From the perspective of the program evaluation literature, this can be viewed as a violation of the Stable Unit Treatment Value Assumption (SUTVA), and in particular a violation of the no-interference assumption: the outcome (the price) of the policy at an "untreated" housing unit in the market is affected by the fact that other housing units were "treated" with changes to their amenities.

To analyze such effects, we can draw on extensions to the potential outcome framework made by Hudgens and Halloran (2008), Tchetgen Tchetgen and VanderWeele (2012), VanderWeele and Tchetgen Tchetgen (2011), and others to consider effects defined by an entire policy—that is, a change in g anywhere. In particular, we can distinguish between the "direct," the "indirect", and the "total" effects of a change in amenities. Consider for simplicity a binary treatment, which occurs only in the second period: $g^0=0, g^1 \in \{0,1\}$. Examples might include cleanup of Superfund sites (Gamper-Rabindran and Timmins 2013, Greenstone and Gallagher 2008), discovery of a cancer cluster (Davis 2004), closing of large polluting facilities (Currie et al. 2015, Davis 2011), and so forth. The amenity itself might be taken to be binary (presence or absence of a Superfund site, for example), but this is not necessary. Conceptually, one might also consider a continuous measure of the amenity in the baseline, but still model the average effect of a binary treatment. For example, one could condition on a baseline measure of toxicity of hazardous waste sites and still estimate the average effect of a cleanup program. In that case,

the baseline toxicity level might be embedded in \mathbf{x} .

Let $g_h^{1,a}$ be the value of g at house h which is realized under some potential scenario a at $t=1$ and let $\mathbf{g}_{-h}^{1,a}$ be the $(H-1)$ -dimensional vector of g at all houses *except* h in scenario a . Let $\mathbf{x}_h^{1,a}(g_h^{1,a})$ be the value of \mathbf{x} at house h in scenario a , which itself is a function of $g_h^{1,a}$. Let a^* be the scenario that was actually implemented, such as a program to clean up toxic waste. Likewise, let a' be some alternative counterfactual scenario that could have prevailed at $t=1$, the outcomes under which one wants to compare to the outcomes under a^* . One natural choice for a' might be "no cleanup program," which would mean $g_h^{1,a'}=0$ for all h . With this notation, different scenarios a describe different possible distributions of g in $t=1$. Although it may appear redundant, we keep the $t=1$ superscript as well to highlight the fact that there may be other changes in the economic environment between $t=0$ and $t=1$.

We incorporate the violation of SUTVA by allowing for different potential prices at house h based not only on the value of g_h , but also based on the entire policy vector \mathbf{g} . The potential outcome for house h , if the rest of the market were under policy a^* , can be written as $p_h^1(\mathbf{g}_{-h}^{1,a^*}, g_h^1 = 1, \mathbf{x}_h^{1,a^*}(g_h^1 = 1))$ if house h were treated and as $p_h^1(\mathbf{g}_{-h}^{1,a^*}, g_h^1 = 0, \mathbf{x}_h^{1,a^*}(g_h^1 = 0))$ if house h were not treated. The potential outcome for house h , if it were not treated and we were in the counterfactual state, would be $p_h^1(\mathbf{g}_{-h}^{1,a'}, g_h^1 = 0, \mathbf{x}_h^{1,a'}(g_h^1 = 0))$. Despite the fact that house h is untreated in either case, $p_h^1(\mathbf{g}_{-h}^{1,a^*}, g_h^1 = 0, \mathbf{x}_h^{1,a^*}(g_h^1 = 0))$ is not necessarily equal to $p_h^1(\mathbf{g}_{-h}^{1,a'}, g_h^1 = 0, \mathbf{x}_h^{1,a'}(g_h^1 = 0))$ because the treatments at the other houses effect equilibrium prices at all houses, including h . Also, even if a' is "no program," $p_h^1(\mathbf{g}_{-h}^{1,a'}, g_h^1 = 0, \mathbf{x}_h^{1,a'}(g_h^1 = 0))$ is not necessarily equal to p_h^0 , as there could be other changes in the economic environment between $t=0$ and $t=1$ affecting the price function or \mathbf{x}_h . Recall any such changes are reflected in the time superscript of the price function $p_h^t(\cdot)$.

Following Hudgens and Halloran (2008), define the individual *total effect* (TE) of treatment by program a^* for some house h as

$$TE_h(a^*) = p_h^1(\mathbf{g}_{-h}^{1,a^*}, g_h^1 = 1, \mathbf{x}_h^{1,a^*}(g_h^1 = 1)) - p_h^1(\mathbf{g}_{-h}^{1,a'}, g_h^1 = 0, \mathbf{x}_h^{1,a'}(g_h^1 = 0)).$$

The total effect is the overall effect of the treatment by the policy at house h . It can be decom-

posed into two parts, an *indirect effect* and a *direct effect*. The individual indirect effect (IE) of treatment a^* for h is defined as

$$IE_h(a^*) = p_h^1(\mathbf{g}_{-h}^{1,a^*}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^{1,a^*}(\mathbf{g}_h^1 = 0)) - p_h^1(\mathbf{g}_{-h}^{1,a'}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^{1,a'}(\mathbf{g}_h^1 = 0)).$$

IE represents the effect on the price of untreated houses due to the shifting hedonic price function between scenario a' and scenario a^* . It is the result of interference: the price at h may be affected by spillovers from treatments elsewhere, even if h itself is untreated.

Following Hudgens and Halloran (2008) and VanderWeele and Tchetgen Tchetgen (2011), define the individual *direct effect* (DE) of treatment a^* for h , conditional on the program going forward in the rest of the market, as

$$DE_h(a^*) = p_h^1(\mathbf{g}_{-h}^{1,a^*}, \mathbf{g}_h^1 = 1, \mathbf{x}_h^{1,a^*}(\mathbf{g}_h^1 = 1)) - p_h^1(\mathbf{g}_{-h}^{1,a^*}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^{1,a^*}(\mathbf{g}_h^1 = 0)).$$

DE_h represents the effect of moving house h from an untreated to a treated state, while holding constant the treatment program at the other houses.

TE, IE, and DE all include any effects mediated through changes in \mathbf{x} .¹² For example, improvements in public goods might attract gentrifying households who improve the house in other (observable) ways. Variants of these treatment effects that net out the portions mediated through changes in \mathbf{x} can be defined for all three. Define the *total unmediated effect* (TUE), the *indirect unmediated effect* (IUE), and the *direct unmediated effect* (DUE) at $\tilde{\mathbf{x}}_h$ respectively by:

$$TUE_h(a^*) = p_h^1(\mathbf{g}_{-h}^{1,a^*}, \mathbf{g}_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) - p_h^1(\mathbf{g}_{-h}^{1,a'}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h).$$

$$IUE_h(a^*) = p_h^1(\mathbf{g}_{-h}^{1,a^*}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) - p_h^1(\mathbf{g}_{-h}^{1,a'}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h).$$

$$DUE_h(a^*)| = p_h^1(\mathbf{g}_{-h}^{1,a^*}, \mathbf{g}_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) - p_h^1(\mathbf{g}_{-h}^{1,a^*}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h).$$

DE and DUE both can be identified by difference-in-differences hedonics. But DUE is the key causal effect concept which has a welfare interpretation as a lower bound on ES. Accordingly, henceforth I shall focus on DUE. It represents moving h from an untreated to a treated state,

¹² As noted by Tchetgen Tchetgen and VanderWeele (2012), terms like "direct" and "indirect" can be somewhat misleading in the presence of both interference and mediation. I use "indirect" to mean changes in the hedonic price function (as defined previously) and "mediated" to be the effect through changes in \mathbf{x} attributable to changes in \mathbf{g} .

while holding constant the treatment program at the other houses and holding constant \mathbf{x}_h at some level $\tilde{\mathbf{x}}_h$.

In principle, changing the treatment status of just h , as envisioned in the definitions of DE and DUE, could itself have general equilibrium effects. The following assumption rules such effects out for DUE.

ASSUMPTION A6 (Local Non-interference). Let $\mathbf{p}_{-h}^1(\mathbf{g}_{-h}^{1,a}, \mathbf{g}_h^1)$ be the vector of prices for all houses except h given a particular treatment status of house h and all other houses. Assume $\mathbf{p}_{-h}^1(\mathbf{g}_{-h}^{1,a}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) \approx \mathbf{p}_{-h}^1(\mathbf{g}_{-h}^{1,a}, \mathbf{g}_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h)$ for all h and all a .

In other words, changing the treatment status of only one house does not appreciably affect the price of any other house. The local non-interference assumption can be taken to be a minimal instance of Palmquist's (1992) localized externality. Under this assumption, the direct effects (DE and DUE) can be interpreted as the movement along a constant hedonic price function, specifically the one prevailing in scenario a^* , from an untreated to a treated state, at a fixed \mathbf{x} .¹³ With the local non-interference assumption, we can write the potential unmediated effects

$$\begin{aligned} p_h^1(\mathbf{g}_{-h}^{1,a}, \mathbf{g}_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) &= p_h^{1,a}(\mathbf{g}_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h), \\ p_h^1(\mathbf{g}_{-h}^{1,a}, \mathbf{g}_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) &= p_h^{1,a}(\mathbf{g}_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h), \end{aligned}$$

as each price function no longer depends on whether h alone was actually treated. That is, the evaluation of $p_h^{1,a}(\mathbf{g}_h^1, \mathbf{x}_h^1)$ would not have to account for the general equilibrium price effects of changing the treatment status of only house h .

With this additional assumption, the direct and indirect effects are depicted in Figure 3a, which shows the hedonic function rather than its first derivative. Absent program a^* , a house would be priced at point P_A in the figure on the counterfactual price function $p^{1,a'}(\cdot)$. If it were not treated but the policy went forward, its price would be P_B on the p_h^{1,a^*} hedonic price function.

¹³ To avoid the general equilibrium effects of changing the treatment for one unit, Hudgens and Halloran (2008) impose a particular randomization assumption that under any policy a , the number of treated units is fixed. Thus, \mathbf{g}_{-h}^{1,a^*} is actually conditioned on the value of \mathbf{g}_h^1 in their definition of the direct effect. To avoid this awkward construction, VanderWeele and Tchetgen (2011) propose an alternative definition of the direct effect in which \mathbf{g}_{-h}^{1,a^*} is fixed, but which no longer decomposes the total effect. The local non-interference assumption provides an alternative way to address this issue. Under this assumption, both definitions of the direct effect are equivalent.

If it were treated its price would be P_C . The total effect is $P_C - P_A$; the indirect effect is $P_B - P_A$; the direct effect is $P_C - P_B$.

Each of the individual effects defined above have their respective group averages. In particular, define the average total unmediated effect and the average direct unmediated effect as follows.

$$\overline{TUE(a^*)} = \frac{1}{H} \sum_h \left[p_h^{1,a^*}(g_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) - p_h^{1,a'}(g_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) \right].$$

$$\overline{DUE(a^*)} = \frac{1}{H} \sum_h \left[p_h^{1,a^*}(g_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) - p_h^{1,a^*}(g_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) \right].$$

Finally, we define the average direct unmediated effect on the treated as

$$\overline{DUET(a^*)} = \frac{1}{\sum_h g_h^{a^*}} \sum_h \left[p_h^{1,a^*}(g_h^1 = 1, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) - p_h^{1,a^*}(g_h^1 = 0, \mathbf{x}_h^1 = \tilde{\mathbf{x}}_h) \right] g_h^{a^*}.$$

Other average treatment effects for other effects or averaged over other domains could be defined similarly.

The recent hedonic literature seems to have settled on the average total unmediated effect, $\overline{TUE(a^*)}$, as the meaning of a "capitalization effect." It captures both the treatment on h and the shifting hedonic price function. If we wanted to forecast the effects on prices of the program, relative to a counterfactual of no program, either $\overline{TE(a^*)}$ or $\overline{TUE(a^*)}$ would be useful measures. However, the impact on prices *qua* prices are of limited economic interest (except of course to the individuals who pay them or receive them as income!). As Kuminoff and Pope (2014) and Klaiber and Smith (2013) have rightly emphasized, $\overline{TUE(a^*)}$ is not the average willingness to pay for program a^* . It conflates the direct and the indirect unmediated effects. Indeed, it is hard to give it any welfare interpretation except in the special case where the hedonic function does not in fact change, in which case the results of Palmquist (1992) apply.

But $\overline{TUE(a^*)}$ is just a straw man. Scenario a' never actually happens and $p_h^{1,a'}$ is never observed for any units, so without additional assumptions this total effect cannot be identified anyway.¹⁴ More to the point, it is not what is identified in most difference-in-differences studies.

¹⁴ That is not to say that, if it were of interest, TE could not be identified with additional assumptions or data. One possible assumption is that there are no other changes in the economic environment, so that, if

In fact, as I show in the following sub-sections, under standard difference-in-differences assumptions, it is the direct effects that can be identified and the direct unmediated effects (\overline{DUE} or \overline{DUET}) that have a clear welfare interpretation.

The basic argument can be seen in Figure 3b. The figure is the same as Figure 3a except the top hedonic function is now the observed ex ante scenario instead of the hypothetical counterfactual scenario a' . Absent the increase in g , both a treated house and its matched counterfactual effect would have begun at P_A . The treated house moves to P_C and the control moves to P_B , at the original level of g but still on the new hedonic price function. Thus, the identified effect from a difference-in-differences comparison is $(P_C - P_A) - (P_B - P_A) = (P_C - P_B)$, which is the movement along the ex post hedonic price function from treatment. If such difference-in-differences are what we mean by the term "capitalization," then it is the change along an equilibrium price function in the direction of g : $p^{1,a}(g^1, \tilde{\mathbf{x}}) - p^{1,a}(g^0, \tilde{\mathbf{x}})$, not the change across time $p^{1,a^*}(g^{1,a^*}, \tilde{\mathbf{x}}) - p^0(g^0, \tilde{\mathbf{x}})$, nor the change across counterfactual equilibria $p^{1,a^*}(g^{1,a^*}, \tilde{\mathbf{x}}) - p^{1,a'}(g^{1,a'}, \tilde{\mathbf{x}})$.¹⁵

3.3 Identification and estimation of capitalization effects: The linear case

Consistently with the vast majority of hedonic work, let us first develop the argument with a linear model. For any individual house h , the ex ante and a^* hedonic price functions and their difference are, respectively:

$$(15a) \quad p_h^0 = \alpha^0 + \beta^0 \mathbf{g}_h^0 + \gamma^0 \mathbf{x}_h^0 + \xi_h + \varepsilon_h^0,$$

$$(15b) \quad p_h^{1,a^*} = \alpha^{1,a^*} + \beta^{1,a^*} \mathbf{g}_h^{1,a^*} + \gamma^{1,a^*} \mathbf{x}_h^{1,a^*} + \xi_h + \varepsilon_h^{1,a^*},$$

a' were "no policy," then p_h^0 could be substituted for $p_h^1(\mathbf{g}_{-h}^{1,a'}, \mathbf{g}_h^1 = 0)$ in the expression for TE . Another is that observations are available at other markets that are not treated and that between-market trends are assumed to be such that identification can leverage inter-city comparisons. See Hudgens and Halloran (2008) and Manski (2013). Crépon et al. (2013) illustrate the idea.

¹⁵ Arguably, this was an early interpretation of hedonics. When Frederick Waugh (1929) first considered hedonic methods for explaining the effect of quality factors on vegetable prices in his PhD dissertation, Eveline Burns, a professor at Columbia, commented that the effect of quality on prices was analogous to the Ricardian rent for quality-differences in the fertility of land (p. 99). By the 1980s, the term "capitalization" was being used interchangeably in potentially confusing ways to refer to cross sectional capitalization or to capitalization over time (compare Bartik 1988, Kanemoto 1988, Scotchmer 1986, and Starrett 1981).

$$(15c) \quad dp_h^{a*} = d\alpha^{a*} + d\beta^{a*} g_h^0 + \beta^{1,a*} dg_h^{a*} + d\gamma^{a*} \mathbf{x}_h^0 + \gamma^{1,a*} d\mathbf{x}_h^{a*} + d\varepsilon_h^{a*},$$

where the differences are taken from the baseline in (15c), not from the unobserved counterfactual. Note the local non-interference assumption A6 is implicitly embedded in (15b), as the parameters are independent of the value of $g_h^{1,a*}$. If the hedonic price function does not change between (15a) and (15b), as with Palmquist's (1992) localized externality, then we can suppress time-scenario superscripts in the parameters and Equation (15c) collapses to

$$(16) \quad dp_h^{a*} = \beta dg_h^{a*} + \gamma' d\mathbf{x}_h^{a*} + d\varepsilon_h^{a*}.$$

In this case, it is clear that difference-in-differences hedonic regressions identify β , the marginal effect of a change in the attribute, if $d\varepsilon$ is independent of dg after conditioning on $d\mathbf{x}$.

However, in the general case where the hedonic function does shift, either because of the changes in g or because of other changes in the economic environment, the true model is (15c), so (16) of course is mis-specified. In particular, it suffers from omitted variable bias: g^0 and \mathbf{x}^0 belong in the model but are omitted. In their recent discussion of capitalization, Kuminoff and Pope (2014) refer to this problem as "conflation bias," a term they use to underscore the fact that the resulting estimates conflate marginal willingness to pay at a point in time (i.e. β^0 or $\beta^{1,a*}$ in this linear example) with changes in the hedonic price function. As they show, "conflation bias" is an example of omitted variable bias. Clearly, if g^0 and \mathbf{x}^0 are included in the model, as in Equation (15c), the linear model potentially can identify $\beta^{1,a*}$, the ex-post marginal willingness to pay under the scenario. Thus, any flaw in the model arises from failure to properly condition on baseline observables, not with the economic logic of differencing prices from two equilibria *per se*.

Of course, including g^0 and \mathbf{x}^0 in a linear model, as in Equation (15c), may well raise additional estimation issues. In particular, although it allows for the existence of an unobserved time invariant effect, ξ_h , Equations (15) still require a conditional zero mean assumption on $d\varepsilon$ to estimate the full set of parameters in (15c) from OLS. Unfortunately, $d\varepsilon$ may well be correlated with g^0 : for example, houses near high levels of pollution may be depreciating in unobserved ways. However, important (if incomplete) information can still be identified under a weaker conditional independence assumption, in which dg^{a*} is independent of $d\varepsilon^{a*}$ conditional on g^0 and

the other observables: $(d\varepsilon \perp dg \mid g^0, \mathbf{x}^0, d\mathbf{x})$. In particular, such conditional independence would allow identification of β^{1,a^*} , even if the coefficient on g^0 were biased. Thus, in this case, we could still identify β^{1,a^*} , the ex post hedonic coefficient, but not β^0 .

The fact that it is the ex post hedonic price parameter under the realized scenario, β^{1,a^*} , that is identified under the weaker assumptions is the crucial point here. In the context of the linear model, this parameter represents \overline{DUE} or \overline{DUET} (per unit g if g is not binary), the direct effect netting out the mediated effect of any changes in \mathbf{x} . The product $dg_h \cdot \beta^{1,a^*}$ is the movement along the ex post hedonic price function in the dimension of g .

3.4 Estimation of capitalization effects without linearity

The preceding insight extends to nonlinear models as well. Relaxing the linearity inherent in Equations (15), now define the potential outcomes using the following semi-parametric assumptions:

$$(17a) \quad p_h^0 = \boldsymbol{\gamma}' \mathbf{x}_h^0 + \varepsilon_h^0,$$

$$(17b) \quad p^{1,a^*}(g_h^1 = 0) = \boldsymbol{\gamma}_{g^1=0}^{1,a^*}' \mathbf{x}_h^{1,a^*} + \varepsilon_{g^1=0,h}^{1,a^*},$$

$$(17c) \quad p^{1,a^*}(g_h^1 = 1) = \boldsymbol{\gamma}_{g^1=1}^{1,a^*}' \mathbf{x}_h^{1,a^*} + \varepsilon_{g^1=1,h}^{1,a^*}.$$

where the $\boldsymbol{\gamma}$ s include an intercept term. This model again implies the local non-interference assumption A6, as $\boldsymbol{\gamma}_g^{1,a^*}$ does not depend on whether any one house h is actually treated.

In addition, we will require a conditional mean independence assumption on differences in unobservables:

ASSUMPTION A7 (Conditional mean independence in differences):

$$(18) \quad E[\varepsilon_{g^1=0}^{1,a^*} - \varepsilon^0 \mid \mathbf{x}^0, g^{1,a^*} = 1] = E[\varepsilon_{g^1=0}^{1,a^*} - \varepsilon^0 \mid \mathbf{x}^0, g^{1,a^*} = 0].$$

In words, after conditioning on \mathbf{x}^0 , the houses that are actually treated by the policy ($g^{1,a^*}=1$) would have had the same expected value of the trend in unobserved time-varying effects, had they not been treated, as the untreated houses ($g^{1,a^*}=0$).

Under these conditions, as well as the usual requirement of overlapping support, a condi-

tional difference-in-differences estimand can identify the average direct unmediated effect on the treated. This is stated more formally in the following lemma.

LEMMA 3. Given A6, A7, and the model of Equations (17),

$$\begin{aligned}
& E \left[\left(p^{1,a^*} (g^1 = 1) - \boldsymbol{\gamma}_{g^1=0}^{1,a^*} \mathbf{x}^1 \right) - \left(p^0 - \boldsymbol{\gamma}_{g^1=0}^{1,a^*} \mathbf{x}^0 \right) \mid \mathbf{x}^0, g^{1,a^*} = 1 \right] \\
& - E \left[\left(p^{1,a^*} (g^1 = 0) - \boldsymbol{\gamma}_{g^1=0}^{1,a^*} \mathbf{x}^1 \right) - \left(p^0 - \boldsymbol{\gamma}_{g^1=0}^{1,a^*} \mathbf{x}^0 \right) \mid \mathbf{x}^0, g^{1,a^*} = 0 \right] \\
(19) \quad & = E \left[\left(p^{1,a^*} (g^1 = 1) - p^{1,a^*} (g^1 = 0) \right) \mid \mathbf{x}^0, g^{1,a^*} = 1 \right] \\
& - \boldsymbol{\gamma}_{g^1=0}^{1,a^*} \left((d\mathbf{x} \mid \mathbf{x}^0, g^{1,a^*} = 1) - (d\mathbf{x} \mid \mathbf{x}^0, g^{1,a^*} = 0) \right) \\
& = \overline{\text{DUET}(a^*)}.
\end{aligned}$$

Proof: The first equality follows immediately from Equations (17) and Assumption A7. The second equality follows from Equations (17), Assumption A6, and the definition of DUET.

For example, $\overline{\text{DUET}(a^*)}$ might be estimated using the regression-adjusted difference-in-differences matching estimator proposed by Heckman, Ichimura, and Todd (1997). In the linear case, the parameter β^1 represents the marginal contribution of g along the ex post hedonic, holding constant any effects mediated through \mathbf{x} . The estimand defined in Lemma 3 recovers an analogous effect, for those houses actually treated by the policy.

The following sub-section is devoted to the economic interpretation of this estimand. But before turning to that discussion, three comments are in order. First, although this discussion has been within the context of a dichotomous amenity or intervention, the result can be extended to include multi-valued or even continuous treatments along the lines suggested by Imbens (2000) and Hirano and Imbens (2004). See, e.g., Muehlenbachs, Spiller, and Timmins (2012) for a hedonic application.

Second, this discussion has been based on the average direct unmediated effect on the treated, $\overline{\text{DUET}(a^*)}$. The effect on the treated makes most sense for an ex-post welfare evaluation of a policy. However, in principle, one could imagine other economically meaningful questions, such as what the welfare effects might have been for alternative policies that had effected other houses. Setting aside the fact that the hedonic price function might have been different in such a case, one might arguably be interested in treatment effects on other populations. One could

identify $\overline{\text{DUE}}(a^*)$ or a direct unmediated effect on other parts of the sample by appropriately modifying Assumption A7. For example, to identify $\overline{\text{DUE}}(a^*)$, the following additional assumption would also be required:

$$E[\varepsilon^{1,a^*}(g^1 = 1) - \varepsilon^0 | \mathbf{x}^0, g^{1,a^*} = 1] = E[\varepsilon^{1,a^*}(g^1 = 1) - \varepsilon^0 | \mathbf{x}^0, g^{1,a^*} = 0].$$

In other words, after conditioning on \mathbf{x} , the control houses ($g^{1,a^*} = 0$) would have followed the same path as the treated houses ($g^{1,a^*} = 1$) had they been treated.

Third, in some cases one may not want to impose the conditional mean independence assumption A7. Although this assumption is weaker than those required for the standard OLS model, one may still be concerned that changes in unobservables are correlated with the treatment. If so, one could invoke additional exclusion restrictions and use instrumental variables. For example, Chay and Greenstone (2005), considering hedonic regressions of housing prices on air quality, persuasively argue that recessions or local economic shocks can simultaneously reduce housing prices in unobserved ways while improving air quality (by dampening economic activity), thus biasing difference-in-differences (or fixed effects) hedonic estimates of air quality downward. They argue that national ambient air quality thresholds are a plausible source of exogenous variation in air quality. Similarly, Greenstone and Gallagher (2008) argue that a discontinuity in the probability of being assigned to Superfund's National Priority List is an exogenous source of variation in the introduction of the Superfund cleanup program, among communities close to the discontinuity. Thus, they condition on the baseline level of the amenity by comparing communities with similar scores on a toxicity index and consider a natural experiment introducing cleanup within this group. In these ways, instrumental variable or regression discontinuity strategies can enrich this estimation of the entity identified in Lemma 3 or Equation (19), although of course the estimand would be different.

3.5 Welfare interpretation of capitalization effects

The previous sub-section showed that difference-in-difference "capitalization" studies can identify $\overline{\text{DUET}}(a^*)$. $\overline{\text{DUET}}(a^*)$ is a well-defined economic concept. It is the difference along the ex post hedonic price function between the value of a house at the new and old level of the amenity respectively, netting out effects mediated through $d\mathbf{x}$: $p^1(g^1, \mathbf{x}^0) - p^1(g^0, \mathbf{x}^0) =$

$\int_{g^0}^{g^1} \frac{\partial p^1(g, x^0)}{\partial g} dg$. The most important thing to emphasize is that the effect is based on the ex post hedonic function. In other words, the counterfactual $p^{1,a^*}(g^1 = 0)$ is *not* the price in the ex ante scenario; nor is it what the price of the house would have been in the absence of the policy, because that counterfactual equilibrium is never observed. The counterfactual is what the price of the house would have been if its g were not affected by the policy but the policy had otherwise gone forward (and the price function had thus shifted). They are only the same in the case of Palmquist's (1992) localized externality; in general they are different.

What is the economic interpretation of this estimand? The expression $\int_{g^0}^{g^1} \frac{\partial p^1}{\partial g} dg$ is a lower bound on ES for the improvement, i.e. the willingness to accept (WTA) to forego the realized change in g . The argument is quite simple in a partial equilibrium context where there are no supply or implicit price effects on x and no effects on firm profits, so that the only effects are the change in the distribution of g . By a simple revealed preference argument, the household consuming g^1 could save expenditures amounting to $\int_{g^0}^{g^1} \frac{\partial p^1}{\partial g} dg$ by consuming g^0 instead. But because it does not choose to do this, the household's minimum WTA must be greater than this amount.¹⁶ This can be seen immediately in Figure 4, which adds a Hicksian demand curve to Figure 1, $h(g, u^1)$, defined at ex-post utility and representing marginal WTA. The WTA associated with EV would be $\int_{g(p^0, u^1)}^{g^1} h(g, u^1) dg$. In contrast, the WTA associated with ES is $\int_{g^0}^{g^1} h(g, u^1) dg$. The difference is in the point of evaluation. The former is associated with the solutions to the expenditure minimization problem given the two hedonic price functions and u^1 . The latter is associated with the realized change in g for the household. Clearly, $\int_{g^0}^{g^1} \frac{\partial p^1}{\partial g} dg$ is a lower bound on ES = $\int_{g^0}^{g^1} h(g, u^1) dg$. The argument is analogous to the well-known fact that a Paasche quantity index is a lower bound for the value of a quantity change.

¹⁶ This point has recently been emphasized by Griffith and Nesheim (2013). It is more common to point out that $\int_{g^0}^{g^1} \frac{\partial p^0}{\partial g} dg$ is an *upper* bound on compensating surplus, $\int_{g^0}^{g^1} h(g, u^0) dg$. Presumably, more attention has been paid to the ex ante price function because it can be used to estimated bounds on WTP for future shocks to the distribution of g , when the induced change on the hedonic price function is as yet unknown. In our case however, it is the ex post hedonic price function that is identified from the quasi-experimental methods. That entity provides a *lower* bound on *ES*.

In fact, under Assumption A6 (local non-interference) and a variant of Assumption A4 (zero profits), this bound remains true even in general equilibrium with a shift in the entire hedonic price function (with implicit prices changing, in general, for all of the amenities, because of either the policy or other changes in the economic environment), endogenous changes in the supply of \mathbf{x} (for example, with upgrades or additions occurring to the housing stock in response to the policy), and resorting of households. The required variant of Assumption A4 is:

ASSUMPTION A4' (zero profits). The change in profits due to adjustments in \mathbf{x} from their counterfactual level are approximately zero when evaluated at the counterfactual level of g :

$$\int_h [p^{1,a^*}(g_h^{1,a'}, \mathbf{x}_h^{1,a^*}) - p^{1,a^*}(g_h^{1,a'}, \mathbf{x}_h^{1,a'})] dh \approx \int_h [c(g_h^{1,a'}, \mathbf{x}_h^{1,a^*}) - c(g_h^{1,a'}, \mathbf{x}_h^{1,a'})] dh.$$

Replacing $t=0$ with the counterfactual scenario, Assumptions A4 and A4' are the same if the price function and cost functions are additively separable in g and \mathbf{x} or if g does not affect the cost function. Alternatively, A4' could be motivated by the idea that adjustments in \mathbf{x} are equilibrated before the change in g .

We can now state Proposition 5, the key result of this sub-section.

PROPOSITION 5. Given A4' and A6, $\overline{\text{DUET}(a^*)} \leq \text{ES}$ for an exogenous change in the distribution of g , when hedonic prices, households, and landlords adjust to the change endogenously, and when there are other changes in the economic environment potentially shifting the price function or the equilibrium levels of \mathbf{x} . If there are changes in current period demands, then the evaluation is from the ex post perspective.

Proof: See the appendix.

The formal proof in the appendix follows the outline of the verbal argument in Bartik (1988), clarifying a few ambiguous points.¹⁷ Kanemoto (1988) provided a proof of a similar lower bound in a related general equilibrium model of land use; however, his bound is on an unusual variant of CV, not ES.

Thus, in contrast to the suggestion recently made by Kuminoff and Pope (2014) and others, there is a clear welfare interpretation of capitalization effects. However, in contrast to the claim by Greenstone and Gallagher (2008), in general they are not equal to either a Hicksian or a Marshallian measure of a change in g . Rather, difference-in-difference hedonics do identify a

¹⁷ Bartik's argument was for an upper bound on WTP, though the role of Assumption A4' is somewhat vague and it is not clear from the paper whether by WTP he had in mind CV or Hicksian compensating surplus. In fact, the bound is on the Hicksian surplus, not variation, measure. From this upper bound, it is straightforward to show the lower bound on ES.

lower bound on ES for a change in g .

4. Illustrative Simulations

In this section I illustrate the results from the previous two sections by simulating hedonic housing equilibria and shocking the equilibria with changes to g . In brief, 100 economies were simulated, each with 1000 households and 1000 houses. In the base model, households have Cobb-Douglas preferences over g and a numeraire, with unobserved taste parameters on g distributed triangular with nodes (0.1, 0.2, 0.6), and with unobserved income distributed log-normal with mean 11.1 and standard deviation 0.4 (and truncated at \$30,000 and \$180,000). Single crossing holds in this model across income conditional on tastes and across tastes conditional on income, but it does not hold between all pairs of individuals. Consequently, there is no *a priori* reason that the approximation given by Equation (12) will be identical to that given by Equation (10). Below, I also consider an alternative model where tastes vary only by an observable, discrete type. Note \mathbf{x} is omitted from the simulations, which simply saves the need to condition on it in the analysis of the simulated data.

The public good g is uniformly distributed on (1, 3) in the baseline scenario. In the ex post scenario, 50% of observations are "treated" by a policy. The probability of being treated is linearly decreasing over the support of g^1 from 0.75 at $g^1=1$ to 0.25 at $g^1=3$. If a house is treated, its level of g improves such that $g^1 = g^0 + (3-g^0)/3 + 1$. Figure 5 illustrates the levels of g in *one* representative simulation. The top panel shows the level of g in the ex post scenario as a function of its level in the ex ante scenario. The bottom line, along a 45-degree ray, represents untreated houses, whereas the top line represents treated houses. The bottom panel of the figure shows the cumulative distribution functions (CDF) of g in each scenario. It shows the ex ante scenario is uniformly distributed, while the ex post scenario, which 1st-order stochastically dominates it, is not. Equilibrium rents range from 16% to 42% of income, with a mean of 27%, which approximates US expenditure shares for housing.

In the base model, equilibrium prices in each scenario were then perturbed by an error term, normally and independently distributed and calibrated such that the standard deviation of the error was equal to either 1% or 5% of the mean price. This error term can be interpreted as either measurement error in price (the dependent variable) or alternatively as an unobserved characteristic of the home that enters preferences as a perfect substitute for the numeraire good.

Below, I also consider the introduction of an unobserved fixed effect correlated with the public good.

Hedonic price functions were fitted non-parametrically to this noisy data in each scenario with a local quadratic function, with bandwidths separately tuned in each simulation using leave-one-out cross validation. Local derivatives were taken analytically from the estimated quadratic function, and so are locally linear.¹⁸ Figure 6 illustrates the estimated price functions for one representative simulation in which the standard deviation of the error was set to 5% of the mean price. The upper left panel shows the price function fit to the data in the ex ante scenario; the upper right shows the respective relationship in the ex post scenario. The lower left shows the two price functions overlaid. Finally, the lower right shows the derivatives of the two respective price functions with respect to g . While the first three panels suggest the relationship is fairly smooth and convex, the final panel does show that the second derivatives are not constant.

Figure 7 displays the relationship between these estimated derivatives and the households' marginal WTP for the public good. Recall that the first-order conditions for the household imply that these should be equal. The left panels in the figure represent the ex ante scenario and the right panels represent the ex post scenario. The top panels plot the slopes of the hedonic price function alongside the marginal WTP to pay of households occupying those houses. The two tend to run together except at very high levels of g , where the estimated derivatives have difficulty keeping up with the rapid escalation in marginal WTP. The bottom panels plot the estimated slope against true marginal WTP. As suggested by the top panels, the estimated slopes fit marginal WTP well except for those with the highest WTP.

Finally, after computing these price functions and their derivatives, the welfare measures defined by expressions (10) and (12) as well as the lower bound of Section 3 were then computed. Additionally, the exact CV, EV, and ES for each household were calculated.

Table 1 reports the results. The first two columns report the results from the base model described above, one column for each standard deviation of the error term. The first two rows show the "true" welfare measures of average CV, EV, and ES (averaging over households within

¹⁸ The optimal bandwidth for fitting prices was adjusted for fitting the first derivative. I also considered modeling the derivative directly from differenced data and tuning the bandwidth using leave-L-out cross validation as suggested by De Brabanter et al. (2013). However, that approach performed quietly poorly, especially near the endpoints.

a simulation), showing the median and the 5th and 95th percentiles of these averages across the 100 simulations. The median average CV is \$6452 in the first model; the median average EV is \$7249 and the ES is close at \$7306.

The third row in Table 1 shows the value of the Harberger approach when the full information on household sorting is available, as computed by Expression (10) (and expressed as a household average). The median estimate of the average value using this approach is \$6926 in the first model. As expected, this lies between the CV and EV measures. The following row places this estimate on to the unit interval between CV and EV. The median placement is 60.4% of the way from CV to EV, just above the midpoint. Similarly, 90 percent of the observations range between 52.1% and 68.1%. All the estimates lie within the interval bracketed by CV and EV. This illustrates that the "sufficient statistic" approach truly is sufficient in this case. Using only information from the first-stage hedonic price function, we can compute a welfare estimate that approximates non-marginal Hicksian welfare measures.

The fifth row shows the value of the Harberger approach when information on household sorting is not available, but such sorting is "imputed" by assuming households sort in the same order of g across the ex ante and ex post scenarios, as described by Expression (12). This imputation is guaranteed to be correct when households are simply ordered and heterogeneity exhibits the single crossing property. However, as noted above, the single crossing property does not strictly hold in this simulation. Nevertheless, the estimates using Expression (12) are virtually identical to those using Expression (10).

To explore the reasons for this result, Figure 8 illustrates the sorting patterns observed in one illustrative simulation. The first panel shows the marginal WTP for g of each household evaluated at the "average house," against the level of g they actually choose in the ex ante scenario. That is, it shows an index of the household's demand-type for g against its optimal g . If Marshallian bid functions never crossed, this figure would show an increasing function. While it is not strictly increasing, it is nearly so, indicating the any crossings in the bid functions are only very local. The second panel in the figure shows the households level of g in the ex post scenario versus the ex ante scenario. It illustrates that households sort approximately in the same order, though not exactly. Finally, the third panel shows the level of g predicted for a household in the ex ante scenario if it sorted in the same order as in the ex post scenario, against its actual

level of g in the ex ante scenario. It shows that the errors from imposing the single-crossing assumption in this scenario are small and local. Consequently, the hedonic price function is approximate constant over the range of these errors, and the over-all bias is minimal. Thus, even when household sorting is not observed, and even when the single crossing property does not hold locally, approximations based on this property potentially can be quite accurate.

The next row shows the estimate of the lower bound using a treatment effects approach with only ex post data. In particular, I use nearest neighbour propensity score matching to estimate the average direct effect on the treated using the ex post data. I then multiply this value by the share of homes that are actually treated to get the average direct price effect (averaging in the zeroes of untreated houses).¹⁹ The median estimate is \$4687 and the median error as a percent of EV, shown in the following row, is -35.5%. The last two rows show the respective estimates using a difference-in-differences matching approach. In this base model, there are no fixed effects to difference out, so the lower bound estimates are essentially unchanged.

In the second column, I consider the sensitivity of these results to changes in the variance of the additive errors, increasing them by a factor of 5. Though the estimates change somewhat, the pattern of results hold, with the Harberger approach falling between the true CV and EV values and the lower bound yielding an estimate about 36% below ES.

In the next column, I consider an additional model where households differ only by five discrete types, as well as by unobserved continuously distributed income. Whereas in the base model households' tastes for g were distributed triangular with nodes (0.1, 0.2, 0.6), in this model those values are rounded to the nearest value in the set {0.15, 0.25, 0.35, 0.45, 0.55}. The household's type is assumed to be observed, so when sorting is not observed it can now be imputed using Equation (13)—i.e. assuming single crossing within type—rather than Equation (12), which assumed single crossing globally. The Harberger approach performs just as well using single-crossing by type as when household locations are observed in the panel, although this is not surprising given the success of Equation (12) in the first model.

Finally, in a fourth model, I introduce an unobserved fixed component of prices, ξ , which

¹⁹ The estimates shown match the 5 nearest neighbors. Alternative numbers of neighbors and estimates using the average direct effect rather than the average direct effect on the treated yielded similar results. In addition, I considered basing an estimate of the Bartik bound directly from the predicted change in value along the estimated ex post hedonic price function. Again, the results were quite similar.

is correlated both with g^1 and with the treatment conditional on g^1 . The final column of Table 1 reports the results from these simulations. The true CV and EV remain unchanged from the first model. However, as seen in rows 3 and 4, the Harberger estimate no longer approximates CV and EV and indeed is 27.2% to 90.8% higher than EV rather than falling within the range between CV and EV. This bias occurs because of ξ , which is correlated with g , so the cross-sectional models recover biased estimates of $\partial p/\partial g$ in both scenarios. Rows 5 and 6 show that the estimated lower bound, based on only the ex post cross section, is now much higher than the estimates in the first model. Again, because of ξ , the results from this model are a biased estimate of DUET. Yet they remain lower bound measures, as in this model this bias merely offsets part of the downward bias relative to EV inherent in the lower bound measure. Finally the difference-in-difference capitalization measures are shown in rows 7 and 8. These effectively condition on ξ and recover an unbiased measure of DUET, or the movement along the ex-post hedonic price function, which is the Bartik lower bound.

5. Conclusions

For decades, economists have used the hedonic model to estimate demands for the implicit characteristics of differentiated commodities, including the demands for otherwise unpriced local public goods and amenities. The traditional cross-sectional approach to hedonic estimation has recovered marginal willingness to pay for amenities when unobservables are conditionally independent of the amenities, but has faltered over a difficult endogeneity problem when attempting non-marginal welfare measures. In this paper, I show that when marginal prices can be reliably estimated, and when panel data on household sorting is available, one can construct an approximation—using only the first-stage marginal prices—which is a "sufficient statistic" for non-marginal welfare measures. With this approximation, Rosen's second-stage estimation can be replaced with a simple average of first-stage parameters. Moreover, even when panel data on household sorting are unavailable, and only repeated cross sections of housing prices (together with a panel on house characteristics) are available, the sufficient statistic approach remains valid under a single crossing restriction. In practice, this approximation appears to perform well in these simulations even when this restriction does not strictly hold.

However, more recently, economists also have questioned the assumptions under which one can identify these cross-sectional hedonic price functions, raising the possibility of unob-

servables that are correlated with the amenity of interest (e.g. Chay and Greenstone 2005, Currie et al. 2015). They have introduced panel econometric models using difference-in-differences and related approaches to identify capitalization effects. Unfortunately, the interpretation of these effects has not been clearly perceived in the literature. In this paper, I show that these capitalization effects identify what is known in the causal literature as the "average direct effect" on prices of a change in amenities, which in this case can be interpreted as a movement along the ex post hedonic price function. I show that this is a lower bound measure on Hicksian equivalent surplus, as suggested by Bartik (1988). The results of the simulations justify taking this approach, as in the presence of unobservables the cross-sectional models are biased in an unknown direction, while the difference-in-differences capitalization model recovers an unbiased estimate of a lower bound (i.e., the bias relative to the true welfare measure is known).

The logic of this paper suggests additional possibilities, hybrids of the results from Sections 2 and 3. For example, if data are available from three or more periods, one could introduce fixed effects and identify the period-2 and period-3 hedonic price functions. This would allow one to use the "sufficient statistics" approach to approximate a non-marginal Hicksian welfare measure that is still robust to the presence of unobservables.

References

- Abadie, A. 2005. Semiparametric Difference-in-Differences Estimators. *Review of Economic Studies* 72(1): 1-19.
- Athey, S., P. Milgrom, and J. Roberts. 1998. *Robust Comparative Statics*. Draft monograph available at faculty-gsb.stanford.edu/athey/documents/draftmonograph98.pdf.
- Bajari, P., and L. Benkard. 2005. Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach. *Journal of Political Economy* 113(6): 1239-76.
- Banzhaf, H.S. 2010. The Chicago School of Welfare Economics. In *The Elgar Companion to the Chicago School of Economics*, ed. by R. Emmett. Cheltenham, UK: Elgar, pp. 59-69.
- Bartik, T.J. 1987. The Estimation of Demand Parameters in Hedonic Price Models. *Journal of Political Economy* 95(1): 81-8.
- Bartik, T.J. 1988. Measuring the Benefits of Amenity Improvements in Hedonic Price Models. *Land Economics* 64(2): 172-83.
- Bayer, P., M.D. Casey, F. Ferreira, and R. McMillan. 2012. Price Discrimination in the Housing Market. NBER working paper 18069.
- Bayer, P., R. McMillan, and K. Rueben. 2011. An Equilibrium Model of Sorting in an Urban Housing Market. Mimeo, Duke University.
- Bayer, P., F. Ferreira, and R. McMillan. 2007. A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy* 115(4): 588-638.
- Bayer, P., R. McMillan, A. Murphy, and C. Timmins. 2011. A Dynamic Model of Demand for Houses and Neighborhoods. NBER Working Paper 17250.
- Bishop, K.C. 2012. A Dynamic Model of Location Choice and Hedonic Valuation. Mimeo. http://www.public.asu.edu/~kcbisho2/dynamic_migration_fall_2012.pdf.
- Bishop, K.C., and C. Timmins. 2015. Estimating the Marginal Willingness to Pay Function without Instrumental Variables. Mimeo. http://www.public.asu.edu/~kcbisho2/estimating_mwtp_function_without_iv.pdf.
- De Brabanter, K., J. De Brabanter, B. De Moor, and I. Gijbels. 2013. Derivative Estimation with Local Polynomial Fitting. *Journal of Machine Learning Research* 14: 281-301.
- Chay, K.Y., and M. Greenstone. 2005. Does Air Quality Matter? Evidence from the Housing Market. *Journal of Political Economy* 113(2): 376-424.
- Chetty, R. 2009. Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods. *Annual Review of Economics* 1: 451-88.
- Chiappori, P.-A., R.J. McCann, and L.P. Nesheim. 2010. Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness. *Economic Theory* 42(2): 317-54.
- Crépon, B., E. Duflo, M. Gurgand, and P. Zamora. 2013. Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment. *Quarterly Journal of Economics* 128(2): 531-80.

- Currie, J., L. Davis, M. Greenstone, and R. Walker. 2015. Do Housing Prices Reflect Environmental Health Risks? Evidence from More than 1600 Toxic Plant Openings and Closings. *American Economic Review* 105(2): 678-709.
- Davis, L.W. 2004. The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster. *American Economic Review* 94(5): 1693-1704.
- Davis, L.W. 2011. The Effect of Power Plants on Local Housing Values and Rents. *Review of Economics and Statistics* 93(4): 1391-1402.
- Depro, B., and C. Timmins. 2012. Residential Mobility and Ozone Exposure: Challenges for Environmental Justice Policy. In *The Political Economy of Environmental Justice*, ed. by H.S. Banzhaf. Stanford University Press: 115-36.
- Edlin, A.S., and C. Shannon. 1998. Strict Single Crossing and the Strict Spence-Mirrlees Condition: A Comment on Monotone Comparative Statics. *Econometrica* 66: 1417-1425.
- 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.
- Epple, D., M. Peress, and H. Sieg. 2010. Identification and Semiparametric Estimation of Equilibrium Models of Local Jurisdictions. *American Economic Journal: Microeconomics* 2(4): 195-220.
- Epple, D., and H. Sieg. 1999. Estimating Equilibrium Models of Local Jurisdictions. *Journal of Political Economy* 107(4): 645-81.
- Figlio, D.N., and M.E. Lucas. 2004. What's in a Grade? School Report Cards and the Housing Market." *American Economic Review* 94: 591-604.
- Fisher, F.M., and K. Shell. 1972. Taste and Quality Change in the Pure Theory of the True Cost-of-Living Index. Essay I in *The Economic Theory of Price Indices*. New York: Academic Press.
- Gamper-Rabindran, S., and C. Timmins. 2013. Does Hazardous Waste Remediation Raise Housing Values? *Journal of Environmental Economics and Management* 65(3): 345-60.
- Greenstone, M., and J. Gallagher. 2008. Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program. *Quarterly Journal of Economics* 123(3): 951-1003.
- Griffith, R., and L. Nesheim. 2013. Hedonic Methods for Baskets of Goods. *Economics Letters* 120: 284-87.
- Harberger, A. C. 1971. Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay. *Journal of Economic Literature* 9(3): 785-97.
- Heckman, J.J. 2010. Building Bridges Between Structural and Program Evaluation Approaches to Evaluating Policy. *Journal of Economic Literature* 48(2): 356-98.
- Heckman, J.J., H. Ichimura, and P.E. Todd. 1997. Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme. *Review of Economic Studies* 64: 605-54.

- Heckman, J.J., L. Lochner, and C. Taber. 1998. General Equilibrium Treatment Effects: A Study of Tuition Policy. *American Economic Review papers and proceedings* 88: 381-6.
- Heckman, J.J., R.L. Matzkin, and L. Nesheim. 2005. Simulation and Estimation of Hedonic Models. In *Frontiers in Applied General Equilibrium*, ed. by T. Kehoe, T.N. Srinivasan, and J. Whalley. Cambridge: Cambridge University Press.
- Heckman, J.J., R.L. Matzkin, and L. Nesheim. 2010a. Nonparametric Estimation of Nonadditive Hedonic Models. *Econometrica* 78(5): 1569-91.
- Heckman, J.J., R.L. Matzkin, and L. Nesheim. 2010b. Supplement to "Nonparametric Estimation of Nonadditive Hedonic Models": Technical Appendix. https://www.econometricsociety.org/ecta/supmat/6388_extensions.pdf.
- Hirano, K., and G.W. Imbens. 2004. The Propensity Score with Continuous Treatments. In, *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, ed. by A. Gelman and X.-L. Meng. Hoboken, NJ: Wiley, pp. 73-84.
- Hudgens, M.G., and M.E. Halloran. 2008. Toward Causal Inference with Interference. *Journal of the American Statistical Association* 103: 832-42.
- Imbens, G.W. 2000. The Role of the Propensity Score in Estimating Dose-Response Functions. *Biometrika* 87(3): 706-10.
- Kanemoto, Y. 1988. Hedonic Prices and the Benefits of Public Projects. *Econometrica* 56(4): 981-9.
- Keane, M.P., and K.I. Wolpin. 1997. The Career Decisions of Young Men. *Journal of Political Economy* 105(3): 473-522.
- Kennan, J., and J.R. Walker. 2011. The Effect of Expected Income on Individual Migration Decisions. *Econometrica* 79(1): 211-51.
- Klaiber, H.A., and V.K. Smith. 2013. Quasi Experiments, Hedonic Models, and Estimating Trade-offs for Local Amenities. *Land Economics* 89(3): 413-31.
- Kuminoff, N.V. 2012. "Recovering Preferences from a Dual-Market Locational Equilibrium, mimeo.
- Kuminoff, N.V., C.F. Parmeter, and J.C. Pope. 2010. "Which Hedonic Models Can We Trust for to Recover Marginal Willingness to Pay for Environmental Amenities?" *Journal of Environmental Economics and Management* 60(1): 145-60.
- Kuminoff, N.V., and J.C. Pope. 2012. A Novel Approach to Identifying Hedonic Demand Parameters. *Economics Letters* 116(3):374-6.
- Kuminoff, N.V., and J.C. Pope. 2014. Do "Capitalization Effects" for Public Goods Reveal the Public's Willingness to Pay? *International Economic Review* 55(4): 1227-50.
- Kuminoff, N.V., V.K. Smith, and C. Timmins. 2013. The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets. *Journal of Economic Literature* 51(4): 1007-62.
- Lavaine, E. 2014. The Capacity of Hedonic Pricing Analysis to Reflect the Economic Benefits of Air Quality. Mimeo.

- Linden, L., and J.E. Rockoff. 2008. Estimates of the Impact of Crime Risk on Property Values from Megan's Laws. *American Economic Review* 98: 1103-27.
- Manski, C.F. 2013. Identification of Treatment Response with Social Interactions. *Econometrics Journal* 16: S1-S23.
- Milgrom, P., and C. Shannon. 1994. Monotone Comparative Statics. *Econometrica* 62: 157-80.
- Muehlenbachs, L., E. Spiller, and C. Timmins. 2012. Shale Gas Development and Property Values: Differences across Drinking Water Sources. NBER Working Paper 18390.
- Palmquist, R.B. 1992. Valuing Localized Externalities. *Journal of Urban Economics* 31(1): 59-68.
- Palmquist, R.B. 2005a. Property Value Models. In *Handbook of Environmental Economics*, ed. by K.G. Mäler and J.R. Vincent. Amsterdam: North-Holland.
- Palmquist, R.B. 2005b. Weak Complementarity, Path Independence, and the Intuition of the Willig Condition. *Journal of Environmental Economics and Management* 49(1): 103-115.
- Parmeter, C.F., and J.C. Pope. 2013. Quasi-Experiments and Hedonic Property Value Methods. In *Handbook on Experimental Economics and the Environment*, ed. by J.A. List and M.K. Price. Cheltenham, UK: Edward Elgar: 3-66.
- Pollak, Robert A. 1989. *The Theory of the Cost-of-Living Index*. New York: Oxford University Press.
- Rosen, S. 1974. Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy* 82(1): 34-55.
- Scotchmer, S. 1985. Hedonic Prices and Cost/Benefit Analysis. *Journal of Economic Theory* 37(1): 55-75.
- Scotchmer, S. 1986. The Short-Run and Long-Run Benefits of Environmental Improvement. *Journal of Public Economics* 30(1): 61-81.
- Sieg, H., V.K. Smith, H.S. Banzhaf, and R. Walsh. 2004. Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods. *International Economic Review* 45(4): 1047-77.
- Smith, V.K., and H.S. Banzhaf. 2004. A Diagrammatic Exposition of Weak Complementarity and the Willig Condition. *American Journal of Agricultural Economics* 86(2): 455-66.
- Stapleford, T.A. 2011. Aftershocks from a Revolution: Ordinal Utility and Cost-of-Living Indexes. *Journal of the History of Economic Thought* 33(2): 187-222.
- Starrett, D.A. 1981. Land value capitalization in local public finance. *Journal of Political Economy* 89(2): 306-327.
- Tchetgen Tchetgen, E.J., and T.J. VanderWeele. 2012. On Causal Inference in the Presence of Interference. *Statistical Methods in Medical Research* 21: 55-75.
- VanderWeele, T.J., and E.J. Tchetgen Tchetgen. 2011. Effect Partitioning under Interference in Two-Stage Randomized Vaccine Trials. *Statistics and Probability Letters* 81: 861-9.
- Waugh, F.V. 1929. *Quality as a Determinant of Vegetable Prices*. New York: Columbia University Press.

Wilson, R.B. 1993. *Nonlinear Pricing*. New York: Oxford.

Figure 1 Willingness to pay for non-marginal changes in public goods

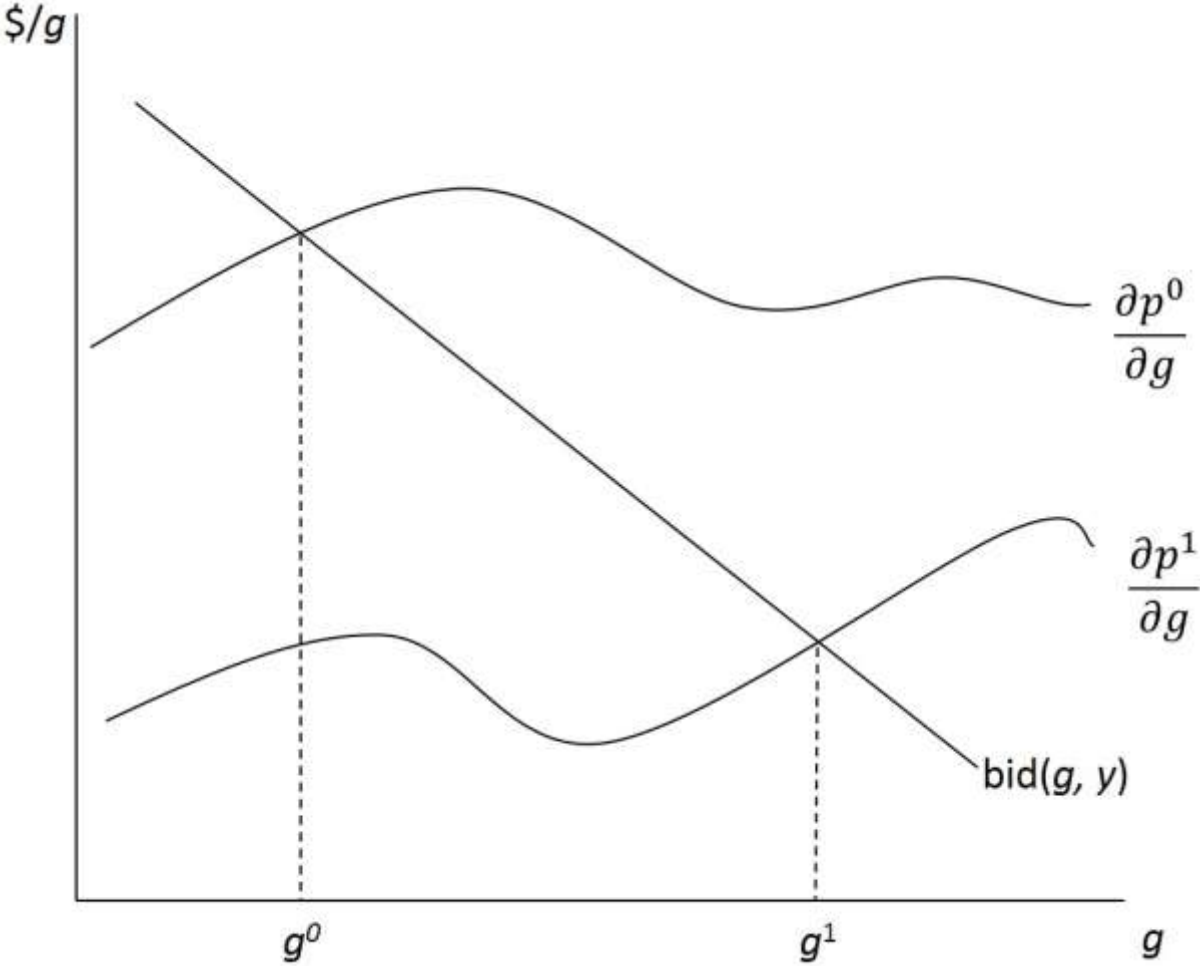


Figure 2 Household sorting is unobserved




Figure 3a. Defining capitalization effects for public goods




Figure 3b. Difference-in-Differences identifies the direct effect




Figure 4. Bounds for WTP and WTA for changes in public goods

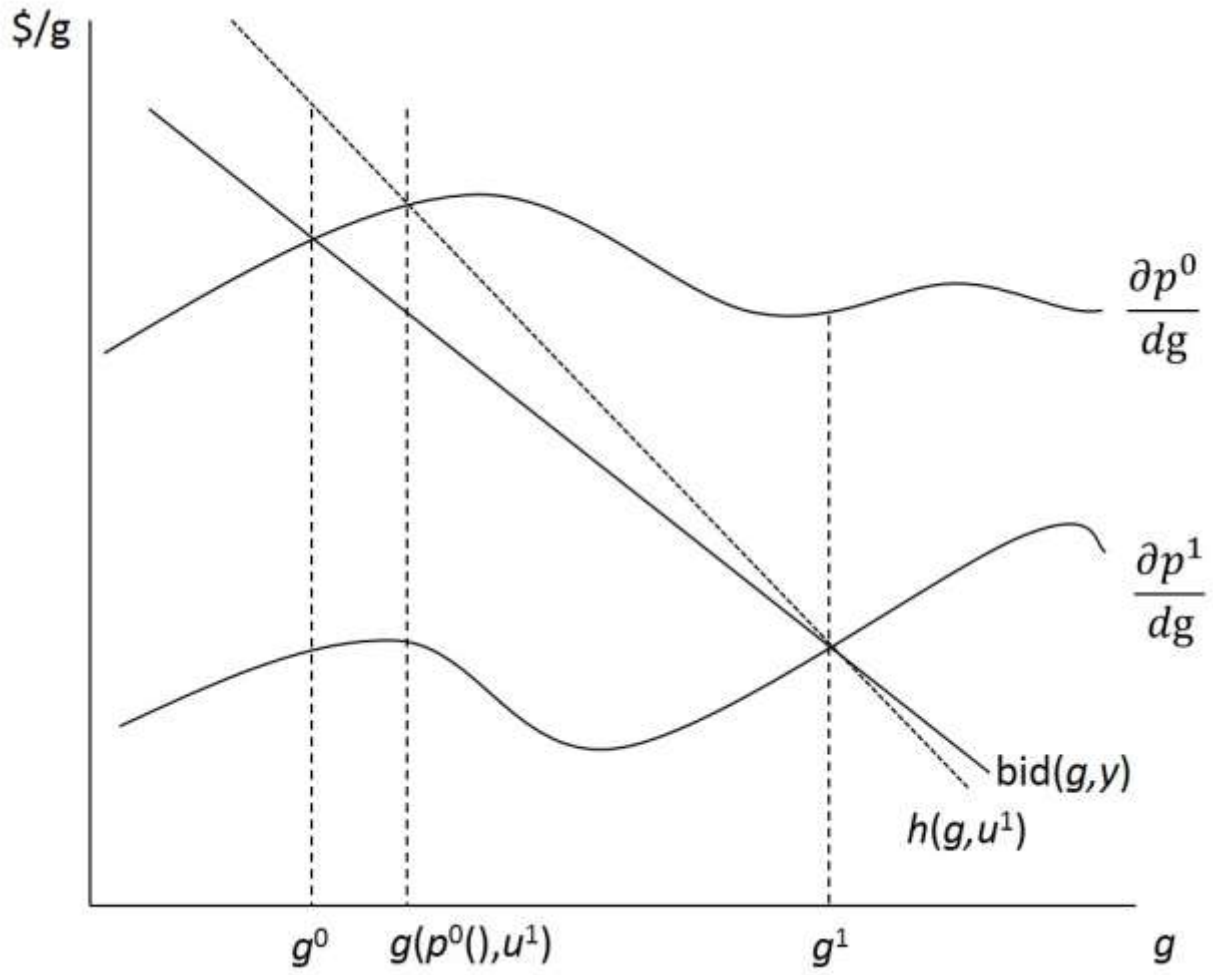



Figure 5. Depiction of Distributions of Public Good for One Representative Policy

5A. Ex post public good vs. ex ante public good (each circle is one house).



5B. CDF of ex post public good (solid) and ex ante public good (dashed)




Figure 6. Depiction of Results for One Representative Simulation: Price Fitting

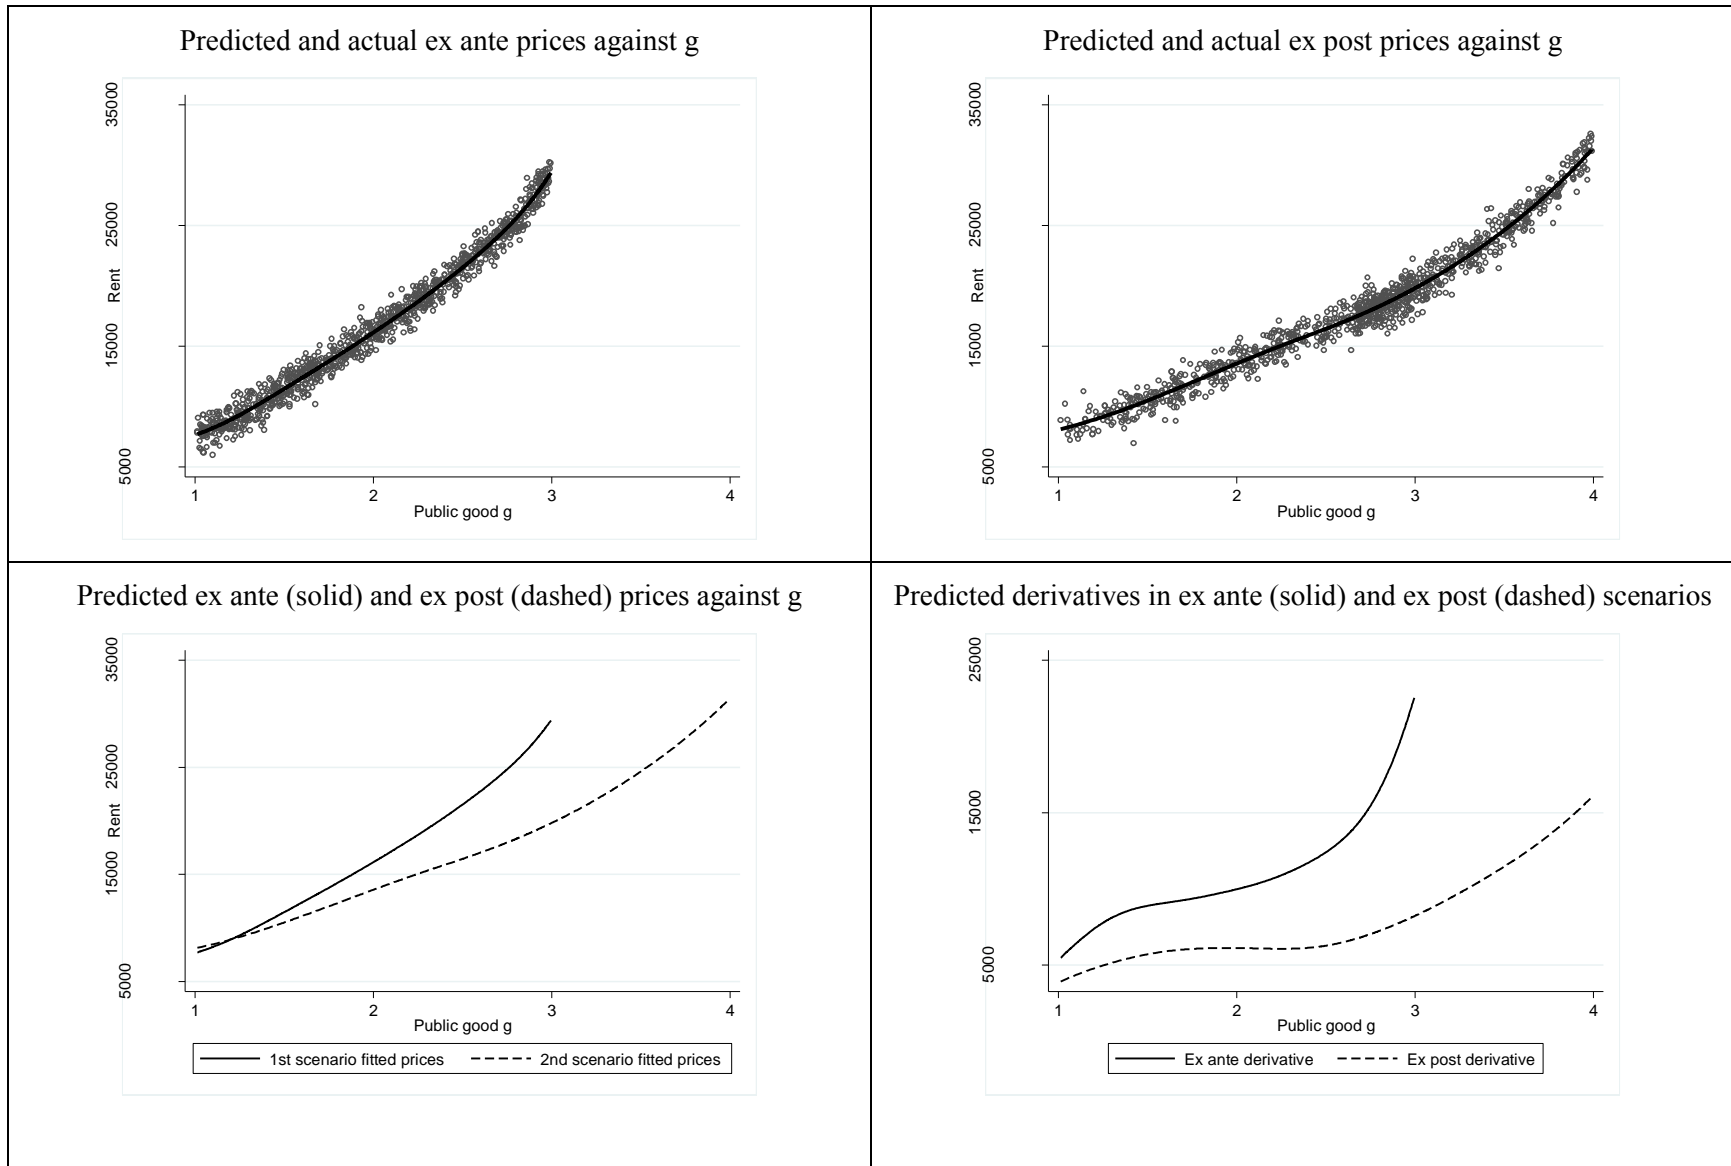


Figure 7. Depiction of Results for One Representative Simulation: Estimated Hedonic Derivative and Marginal WTP

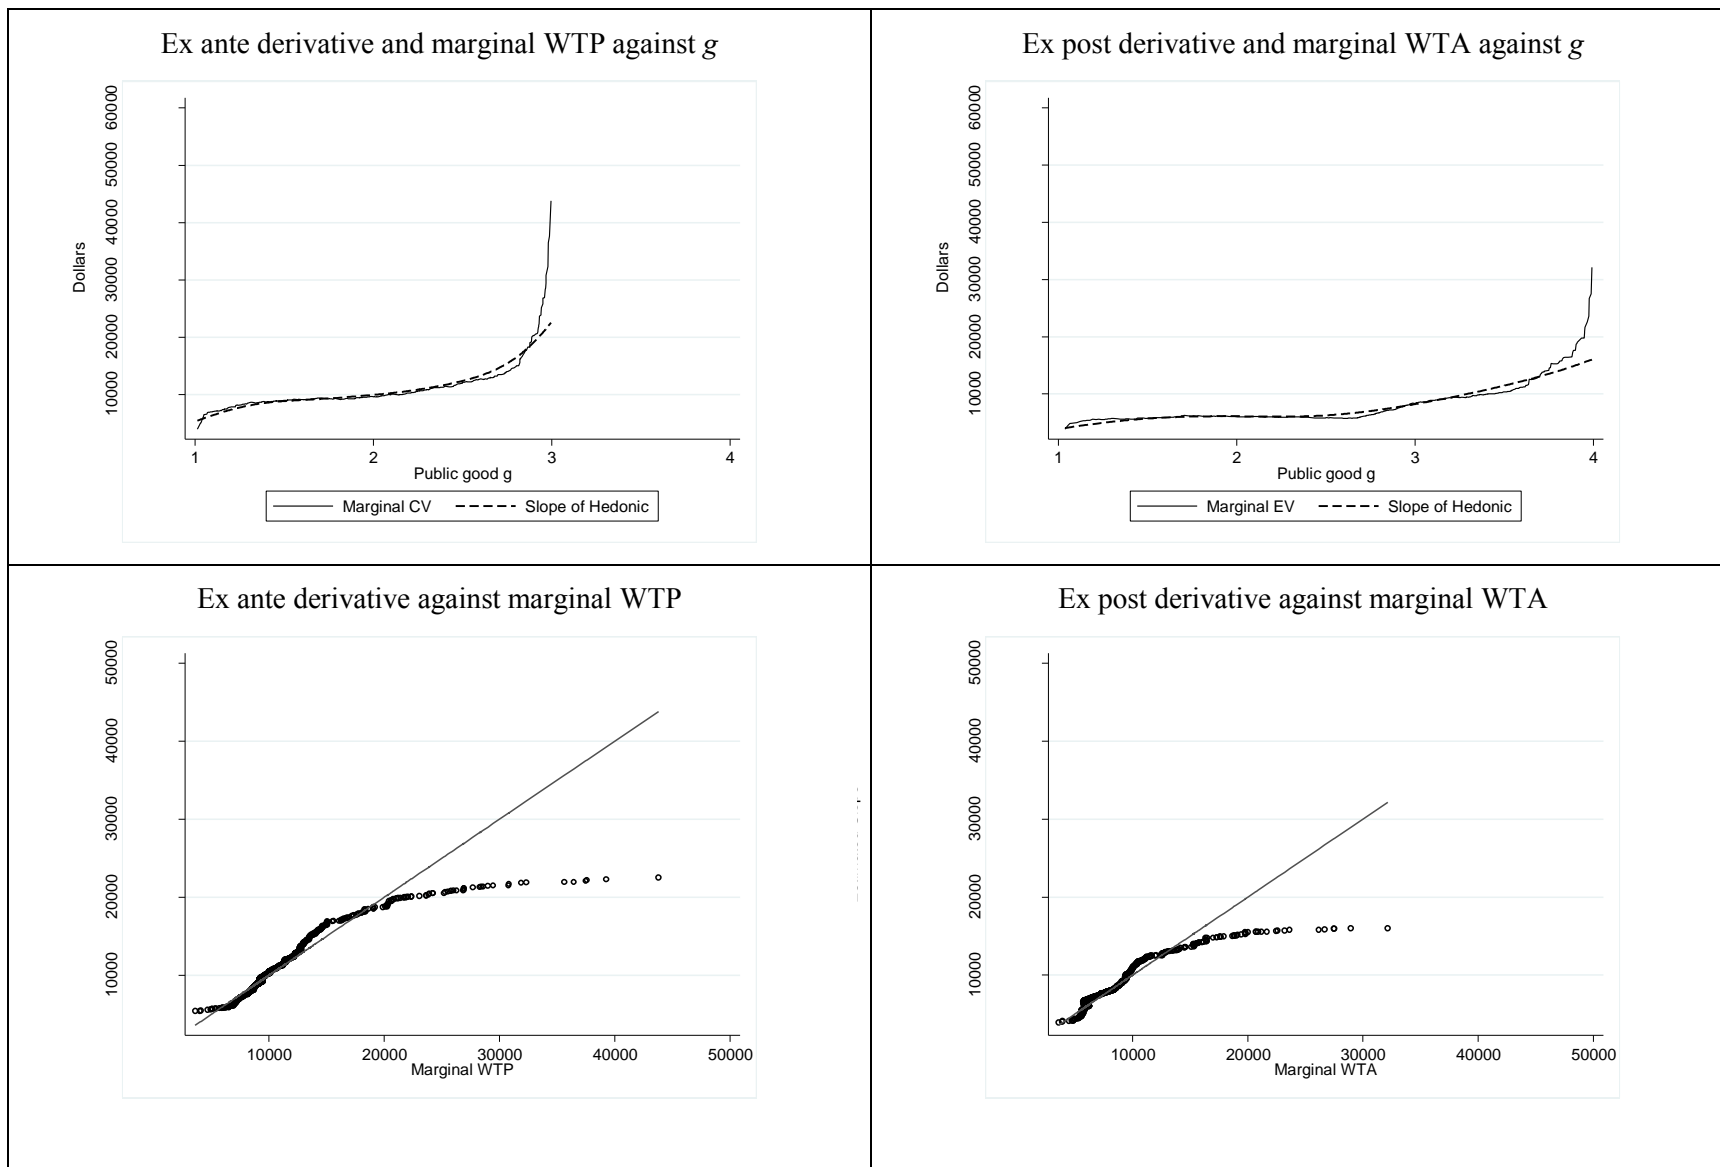


Figure 8. Sorting Behavior




Figure 9. Results for One Simulation with Unobserved Characteristics: Estimated Derivative and Marginal WTP

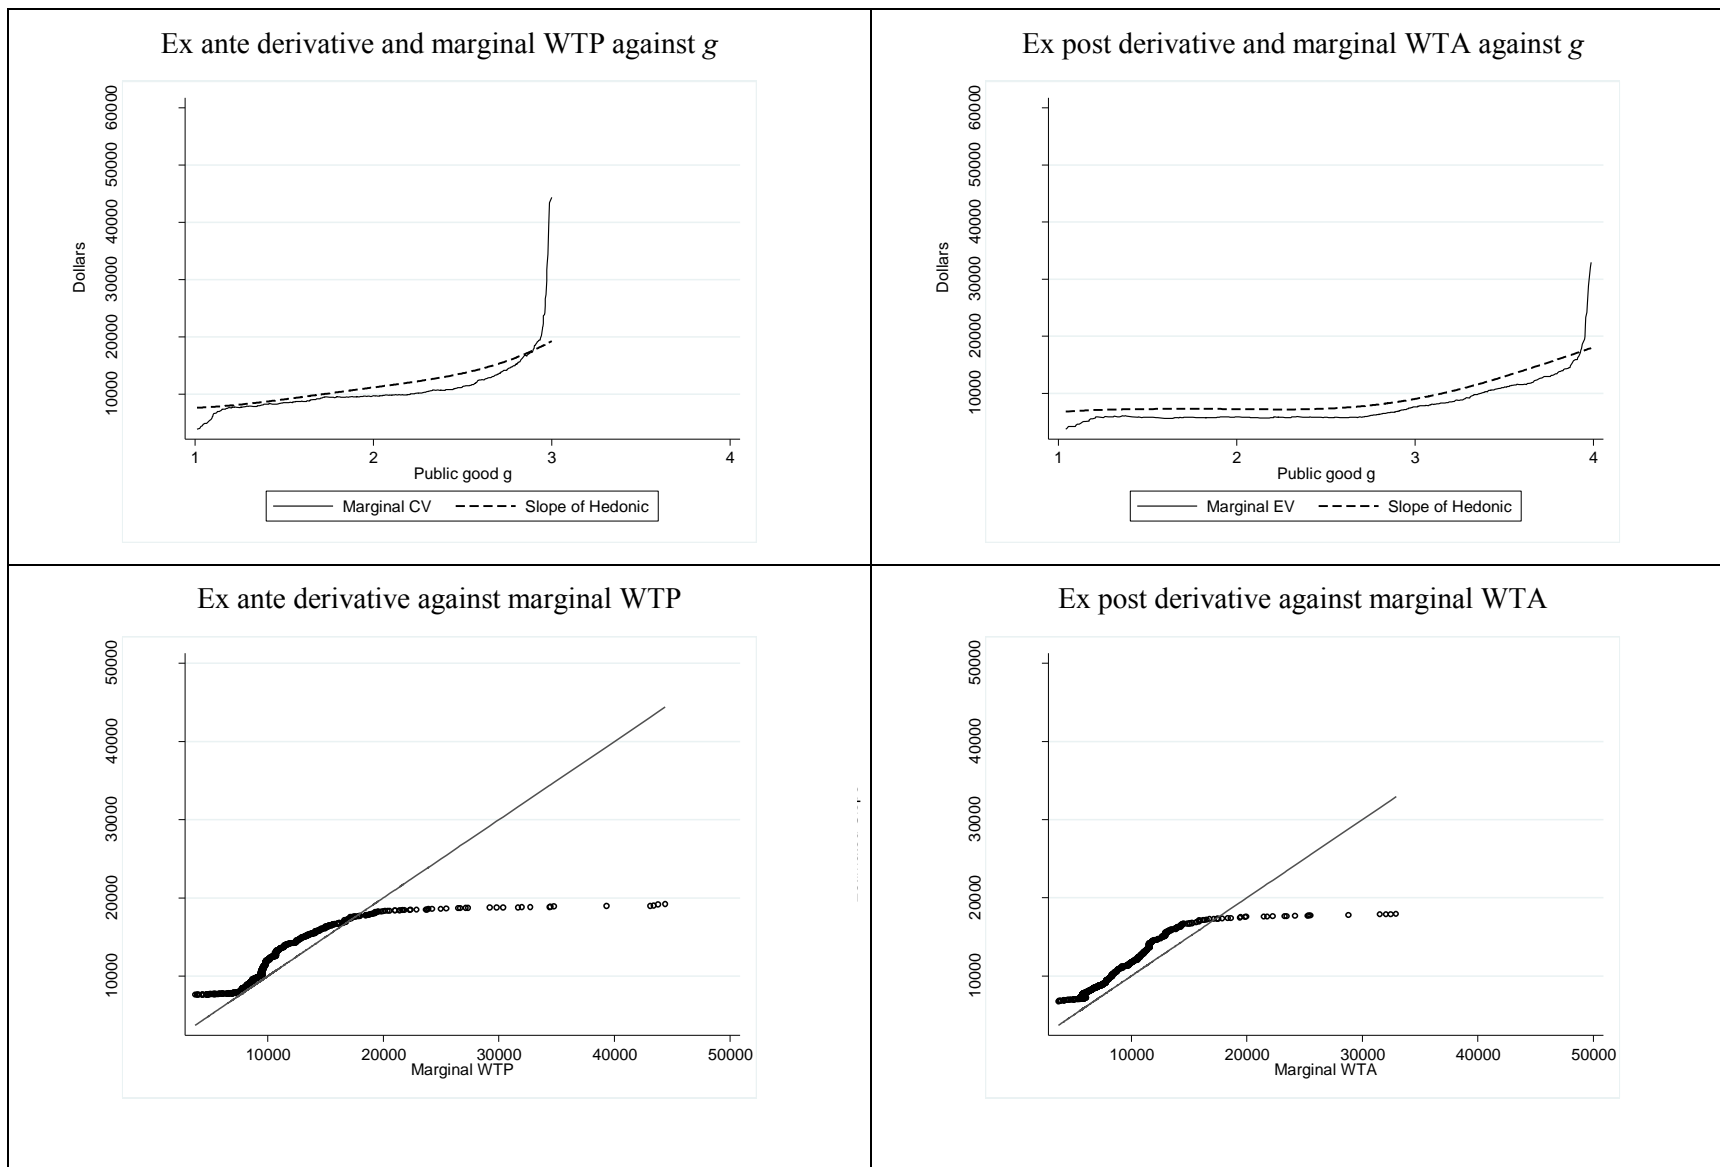


Table 1. Results of Simulations

Statistic (Averaged over households)	Median Value Across 100 Simulations (5 th and 95 th Percentiles in Parentheses)			
	g only (no ξ), type unobserved		g-only (no ξ), type observed	g and ξ , type unobserved
	$\sigma_\varepsilon = 1\%$ of mean price	$\sigma_\varepsilon = 5\%$ of mean price	$\sigma_\varepsilon = 5\%$ of mean price	$\sigma_\varepsilon = 5\%$ of mean price
1. Average CV	\$6452 (\$6044 - \$6787)	\$6452 (\$6044 - \$6787)	\$6518 (\$6140 - \$6825)	\$6452 (\$6044 - \$6787)
2. Average EV	\$7249 (\$6767 - \$7703)	\$7249 (\$6767 - \$7703)	\$7372 (\$6901 - \$7773)	\$7249 (\$6767 - \$7703)
3. Average ES	\$7306 (\$6812 - \$7765)	\$7306 (\$6812 - \$7765)	\$7432 (\$6951 - \$7838)	\$7306 (\$6812 - \$7765)
4. Value of Eq. (10), as avg.	\$6926 (\$6482 - \$7333)	\$6843 (\$6414 - \$7265)	\$6924 (\$6517 - \$7338)	\$7674 (\$7181 - \$8106)
5. Eq.(10) as pct. distance from CV to EV	60.4% (52.1% - 68.1%)	51.6% (35.2% - 67.0%)	50.1% (33.5% - 65.0%)	154.1% (127.2% - 190.8%)
6. Value of Eq. (12) [Eq. (13) when type observed], as avg.	\$6926 (\$6482 - \$7333)	\$6843 (\$6414 - \$7265)	\$6924 (\$6517 - \$7338)	\$7674 (\$7181 - \$8106)
7. Bartik lower bound	\$4687 (\$4448 - \$4869)	\$4693 (\$4441 - \$4862)	\$4690 (\$4467 - \$4884)	\$5374 (\$5087 - \$5598)
8. Avg. Pct. error in Bartik lower bound relative to ES	-35.5% (-37.7% - -34.4%)	-35.6% (-37.9% - -34.3%)	-36.6% (-39.0% - -35.2%)	-26.0% (-28.7% - -24.4%)
9. Bartik lower bound using difference-in-differences (DD)	\$4687 (\$4438 - \$4868)	\$4688 (\$4397 - \$4864)	\$4695 (\$4433 - \$4891)	\$4688 (\$4399 - \$4864)
10. Avg. Pct. error in DD Bartik lower bound relative to ES	-35.5% (-37.7% - -34.4%)	-35.6% (-38.1% - -34.3%)	-36.6% (-38.9% - -35.1%)	-35.6% (-38.1% - -34.3%)

APPENDIX. PROOFS OF LEMMAS AND PROPOSITIONS.

Proof of Lemma 1.

Since the marginal utilities are functions of p and \mathbf{z} , we can write Equation (4) in terms of changes in marginal utilities, $d\left(\frac{\partial v_i}{\partial p}\right)$ and $d\left(\frac{\partial v_i}{\partial z_j}\right)$. With this notation, Equation (4) simplifies to:

$$(20) \quad dv_i \approx \frac{\partial v_i}{\partial p} dp_i + \sum_j \frac{\partial v_i}{\partial z_j} dz_{j,i} + \frac{1}{2} \left(d\left(\frac{\partial v_i}{\partial p}\right) dp_i + \sum_j d\left(\frac{\partial v_i}{\partial z_j}\right) dz_{j,i} \right).$$

Taking the total derivative of the household's first-order condition given by (1), we have

$$(21) \quad d\left(\frac{\partial v_i}{\partial z_j}\right) = d\lambda_i \frac{\partial p^0}{\partial z_j} \Big|_{\mathbf{z}_i^0} + \lambda_i^0 d\left(\frac{\partial p}{\partial z_j}\right) + d\lambda_i d\left(\frac{\partial p}{\partial z_j}\right),$$

where $d\left(\frac{\partial p}{\partial z_j}\right) = \frac{\partial p^1}{\partial z_j} \Big|_{\mathbf{z}_i^1} - \frac{\partial p^0}{\partial z_j} \Big|_{\mathbf{z}_i^0}$; that is, it is the change in the derivative of the price function from both the change in the price function itself and from the change in the point where it is evaluated. Inserting this expression along with $\frac{\partial v}{\partial p} = -\lambda$ into (20), we have:

$$(22) \quad \begin{aligned} dv_i \approx & -dp_i \left(\lambda_i^0 + \frac{1}{2} d\lambda_i \right) + \sum_j \left(\frac{\partial p^0}{\partial z_j} \Big|_{\mathbf{z}_i^0} \right) dz_{j,i} \left(\lambda_i^0 + \frac{1}{2} d\lambda_i \right) \\ & + \frac{1}{2} \sum_j d\left(\frac{\partial p}{\partial z_j}\right) dz_{j,i} \lambda_i^0 + \frac{1}{2} \sum_j d\left(\frac{\partial p}{\partial z_j}\right) dz_{j,i} d\lambda_i. \end{aligned}$$

Adding and subtracting $\frac{1}{4} \sum_j d\left(\frac{\partial p}{\partial z_j}\right) dz_{j,i} d\lambda_i$, ignoring the remaining third-order terms (as we are taking a second-order approximation), and re-arranging, we have the desired expression. Note to compute this expression, we need Assumption A3, that the price derivatives and dz terms can be computed for each household across time.

Proof of Lemma 2.

The proof follows a similar outline as that for Lemma 1. Since the marginal costs are functions of \mathbf{x} , we can write Equation (6) in terms of changes in marginal costs $d\left(\frac{\partial c_h}{\partial x_r}\right)$. With this notation, Equation (6) can be re-written as:

$$(23) \quad dc_h \approx \sum_r \left(\frac{\partial c_h}{\partial x_r} + \frac{1}{2} d \left(\frac{\partial c_h}{\partial x_r} \right) \right) dx_{r,h}.$$

Substituting the first-order condition (2), we now have:

$$(24) \quad \begin{aligned} dc_h &\approx \sum_r \left(\frac{\partial p^0}{\partial x_r} \Big|_{\mathbf{x}_h^0} + \frac{1}{2} d \left(\frac{\partial p}{\partial x_r} \right) \right) dx_{r,h} \\ &= \sum_r \frac{1}{2} \left(\frac{\partial p^0}{\partial x_r} \Big|_{\mathbf{x}_h^0} + \frac{\partial p^1}{\partial x_r} \Big|_{\mathbf{x}_h^1} \right) dx_{r,h} \end{aligned}$$

Substituting this expression into the equation $d\pi = dp - dc$ completes the proof.

Proof of Proposition 4.

Consider the initial equilibrium described by the hedonic price function $p^0(g, \mathbf{x}; F^0)$. Now consider a change in the distribution of g and allow preferences for individual households to change, but such that they are still distributed the same (Assumption A2'). Denote the new equilibrium price function by $p^1(g, \mathbf{x}; F^1)$. Now consider a counter-factual scenario where the distribution of g is still given by F^1 but all households have their original demand functions. By Assumption A2', the equilibrium price function would still be $p^1(g, \mathbf{x}; F^1)$ since the distribution of demands would be no different in this counterfactual g and \mathbf{x} are unchanged (though the assignment of households to houses may change). But by Assumption A5 and the argument in Proposition 3, any household choosing z_j in the initial equilibrium such that $F_j^0(z_j) = \theta$ will choose z_j in the counterfactual equilibrium such that $F_j^1(z_j) = \theta$. Moreover, these choices are consistent with the price equilibrium. The rest of the argument follows from Proposition 3.

Finally, note that the same argument could be made, *mutatis mutandis*, starting with $p^1(\cdot)$ and going back to $p^0(\cdot)$ under ex post preferences. Thus, the aggregate welfare evaluation is invariant to the perspective taken.

Proof of Proposition 5.

Denote the expenditure function for household i as $e_i(p(\cdot), u)$ where $p(\cdot)$ is the hedonic price function and the price of other goods is normalized to one. It is the solution to the expenditure minimization problem when the household faces hedonic price function $p(\cdot)$. Denote the restricted expenditure function as $\tilde{e}_i(p(g, \mathbf{x}), g, \mathbf{x}, u)$; it is the solution to the expenditure minimization problem when the household is constrained to choose the bundle (g, \mathbf{x}) .

Denote our measure of the change in welfare by:

$$(25) \quad dW = \sum_i [\tilde{e}_i(p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}), \mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_h [(p^{a*}(\mathbf{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a'}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) - (c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}))].$$

The first term in square brackets is, by definition, the Hicksian equivalent surplus (ES) for the change in g . It differs from Hicksian equivalent variation (EV) insofar as the household is constrained to be at $(\mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'})$ in the first expenditure function, which was the solution to the expenditure minimization problem for $p^{a'}(\cdot)$ at $u_i^{a'}$, not u_i^{a*} . If preferences have changed over time, then this is the current-period ES for $t=1$. The second term in brackets is the change in landlord profits. It is the change in rents, resulting from both the shift in the hedonic price function and adjustments in \mathbf{x} as well as exogenous changes in g , minus the change in costs, evaluated at baseline levels of g .

The right side of Equation (25) can be decomposed as follows:

$$(26) \quad dW = \sum_i [\tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_i [\tilde{e}_i(p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - \tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*})] \\ + \sum_h (p^{a*}(\mathbf{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*})) + \sum_h (p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) \\ + \sum_h (p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}) - p^{a'}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) - \sum_h (c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})).$$

Fixing indices so that $i(a')=h$, which we can do because of the bijective mapping between houses and households, the expression can be re-arranged as

$$(27) \quad dW = \sum_i [\tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_h (p^{a*}(\mathbf{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*})) \\ + \sum_h [(p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) - (c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}))] \\ + \sum_i [(\tilde{e}_i(p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - \tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*})) \\ - (p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}) - p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}))]]$$

The third line is equal to zero by Assumption A4'. Additionally, for each i , the term in the fourth line minus the term in the last line is equal to zero by the definition of \tilde{e} : The money necessary to maintain utility when (g, \mathbf{x}) is held fixed is equal to the change in the price of the bundle

(g, \mathbf{x}). Thus, the entire expression collapses to:

$$(28) \quad dW = \sum_i [\tilde{e}_i(p^{a*}(g_{i(a')}, \mathbf{x}_{i(a')}), g_{i(a')}, \mathbf{x}_{i(a')}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_h (p^{a*}(g_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(g_h^{a'}, \mathbf{x}_h^{a*}))$$

But the term in square brackets is non-negative for each i : the value of a constrained expenditure minimization problem is no less than the value of an unconstrained expenditure minimization problem at the same prices and utility. Thus,

$$(29) \quad \sum_h (p^{a*}(g_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(g_h^{a'}, \mathbf{x}_h^{a*})) \leq dW.$$

This completes the proof. The term on the left is the sum of price changes along the ex-post hedonic holding \mathbf{x} constant at its ex-post level, which is the measurement of interest, and it is less than the welfare measure.