Difference-in-Differences Hedonics

H. Spencer Banzhaf
Georgia State University, NBER, PERC

October 2020

Abstract

Traditional cross-sectional estimates of hedonic price functions theoretically recover marginal willingness to pay for characteristics, but face endogeneity problems. Consequently, economists have introduced difference-in-differences and other quasi-experimental econometric methods into the hedonic model. Unfortunately, the welfare interpretation of these estimands has not been clear. This paper shows that, when they condition on baseline data, quasi-experimental hedonics identify the "movement along" the ex post price function from a shock. It further shows this effect is a lower bound on welfare, one which accounts for non-marginal changes and general equilibrium effects. Thus, information about nonmarginal welfare can be recovered using transparent designs, without Rosen's "second stage" or more structural models. The paper illustrates these results with an application to toxic facilities' effects on housing prices. In this case, the lower bound is up to 20% higher than traditional capitalization approaches.

Difference-in-Differences Hedonics

1. Introduction

For over half a century, economists have used hedonic price functions as a simple way to model quality-differentiated products. The standard hedonic model uses cross-sectional data and the insight that the derivative of the price function with respect to an attribute equals households' marginal willingness to pay (WTP) (Rosen 1974). While working with cross-sections was the standard hedonic project for decades, recently economists have drawn attention to the problem of unobserved attributes that may be correlated with the attribute of interest (Greenstone 2017). To overcome this endogeneity problem, they have applied difference-in-differences (DD) and other quasi-experimental methods to the hedonic econometric model.[1]

While the new approach has definite econometric advantages, it seemingly has come at the cost of clarity about the estimand: What economic question it answers is not always clear, or at least has not been perceived clearly in the literature. For example, if we first-difference the data, the dependent variable becomes a change in prices, which mixes information from two equilibria. Yet households face a tradeoff among products at a point in time, not across time.

In the housing context, the literature has begun to refer to such changes in prices over time as "capitalization" (e.g. Chay and Greenstone 2005, Klaiber and Smith 2013, Kuminoff and Pope 2014). The link between such capitalization and the underlying economic model is not immediately clear. For changes to the characteristics of a small subset of houses, the equilibrium price function can be taken as constant over a short time period, so DD models can be interpreted within a single equilibrium (Palmquist 1992). But in general, a large change in the supply of an amenity will endogenously shift the hedonic price function. Too, other changes in the economic environment taking place over longer time periods also would shift the price function. In these cases, panel-data studies compare prices at two different equilibria in potentially confusing ways. The confusion is

[1] Examples include Bento et al. (2015), Cellini et al. (2010), Chay and Greenstone (2005), Currie et al. (2015), Davis (2004, 2011), Greenstone and Gallagher (2008), Haninger et al. (2017), Lang (2018), Linden and Rockoff (2008), Mastromonaco (2015), Muehlenbachs et al. 2015, and Pope (2008). Parmeter and Pope (2013) review this literature. Kuminoff et al. (2010) illustrate the importance of using DD hedonic techniques to control for unobservables.

compounded by ambiguity about the meaning of language borrowed from the program evaluation literature, such as "*the* capitalization effect," when in fact there are various such effects with differing interpretations.

Klaiber and Smith (2013), Parmeter and Pope (2013), and Kuminoff and Pope (2014) have argued that, because it combines two equilibria, the capitalization effect answers an ill-defined economic question. Taken together, these papers make two key points. First, the total change in prices from an exogenous change in characteristics is not the same as WTP. Instead, it conflates WTP (defined within the context of one equilibrium) with changes in the price function. Second, estimated cross-sectional price functions are biased if they ignore changes in the function. Both points are correct. Nevertheless, as I show in this paper, DD hedonic studies can provide meaningful welfare measures. This paper clarifies the issue by introducing distinctions made in the treatment effect literature when the Stable Unit Treatment Value Assumption (SUTVA) is violated (Section 2). When the hedonic price function shifts endogenously, there are indirect effects on the price of all units, regardless of treatment status, which violates the usual SUTVA condition.

Using this framework, I first show in Section 3 that if they allow for changes in the price function over time, quasi-experimental hedonic studies can identify what this literature calls the "direct effect" of treatment, or the effect of treatment on a given unit, conditional on the treatment program going forward for all other units. This direct effect has the simple interpretation as the movement along the ex post hedonic price function.

Second, in Section 4 I show that this effect can be interpreted as a lower bound on Hicksian equivalent surplus (ES) for a non-marginal change in characteristics of interest, even in the presence of endogenous general equilibrium shifts in the price function, shifts in the price function from other exogenous factors, and endogenous adjustments to other characteristics. The bound is similar to one discussed many years ago by Bartik (1988). These results are quite general. Demands and other aspects of the economic environment may change between periods, there are no restrictions on heterogeneity in demands, household data are not required, and even repeated cross sections can be used. Thus, quasi-experimental estimates of the hedonic price function alone are sufficient to yield at least a bound on non-marginal changes, even accounting for general equilibrium effects, without requiring economists to estimate Rosen's (1974) difficult "second stage" or other more structural models. For truly marginal effects, they collapse to marginal WTP.

Section 5 illustrates these results with an application to the value of reduced toxic emissions in Los Angeles between 1995 and 2000. Although a lower bound on general equilibrium welfare, the estimates are substantial, at about $6.7-7.5 billion for the present value of the realized decrease in emissions, or about $62-69 per household annually. In contrast, the conventional approach to capitalization with a time-invariant function gives values about 14-20% lower. Simply linearizing marginal WTP gives estimates that are lower still. Thus, besides clarifying the interpretation of quasi-experimental hedonics, the method suggested here can tighten the lower bound.

2. Hedonic Capitalization Effects

2.1 The Hedonic Model

Although the results of this paper apply to any hedonic context, for concreteness, consider the specific application of housing price functions. Let $\mathcal{H}$ denote the set of houses in a region with typical element h and let $\mathcal{I}$ denote the set of households with typical element i. For ease of exposition, assume for now that the region is closed, so there is no migration in or out. (The extensions in Appendix A show this assumption can be relaxed without affecting the interpretation.) 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 may rent or own their house, and owners rent from themselves.

At time t, households differ by their income y and their current-period preferences, which are represented by a twice differentiable conditional indirect utility function over prices, an amenity of interest g, and other characteristics $\mathbf{x}$, $v_i^t(y_i^t\text{-}p_h, g_h, \mathbf{x}_h)$, with $\partial v_i^t/\partial y_i^t > 0$. On the supply side, the profit function for house h is $\pi_h = p_h - c_h(\mathbf{x}_h)$, where the cost function $c_h(\ )$ is twice differentiable. For simplicity, assume $c_h(\ )$ is constant over time, although this assumption could be relaxed.

Consider two points in time, denoted t=0 for an initial situation and t=1 in a later situation. In each period, prices of houses are determined by the level of the amenities evaluated on the time-specific equilibrium price function: $p_h^t = p^t(g_h^t, \mathbf{x}_h^t)$. The time superscript on the price function indicates that equilibrium hedonic prices may shift. In each period, households maximize utility over a continuous choice set defined by the continuously differentiable price function.

I make the standard hedonic assumption that households have perfect information and are in a static equilibrium in each time period.[2] Maximizing utility in period *t*, the household satisfies the first-order condition for *g*: $\frac{\partial v_i^t}{\partial g} = -\frac{\partial v_i^t}{\partial p}\frac{\partial p^t}{\partial g}$. Using $-\partial v_i^t/\partial p = \partial v_i^t/\partial y$, this is equivalent to:

$$\frac{\partial v_i^t/\partial g}{\partial v_i^t/\partial y} = \frac{\partial p^t}{\partial g}. \tag{1}$$

Equation (1) is the standard tangency condition from Rosen (1974), with the derivative of the hedonic function with respect to an amenity equal to marginal WTP at the optimal point.

Similarly, a landlord's first-order condition for profit maximization for characteristic x_r is:

$$\frac{\partial c_h}{\partial x_r} = \frac{\partial p^t}{\partial x_r}. \tag{2}$$

The endogenous amenities **x** are supplied according to similar tangency condition, with marginal cost equal to marginal revenue.

The basic problem is (i) to estimate these primitive conditions when *g* or **x** are endogenous and (ii) to make inferences about non-marginal welfare effects from these primitive conditions.

2.2 Defining Capitalization Effects

To overcome the first issue, related to estimation, researchers have applied DD and related quasi-experimental approaches to the hedonic model to identify the effects of exogenous changes in *g*. While effectively addressing an econometric problem, these methods have raised other questions about what precisely they identify (Klaiber and Smith 2013 and Kuminoff and Pope 2014).

If the equilibrium hedonic price function for a housing market changes endogenously because of a shock to amenities, then the price of a house will change even if its amenities have not. In the vocabulary of the program evaluation literature, this is a violation of SUTVA, particularly

[2] This assumption continues to underlie majority of work on hedonic markets (Bajari and Benkard 2005, Bishop and Timmins 2019, Ekeland et al. 2004, Heckman et al. 2010) as well as structural models of locational choice (Epple et al. 2020). More recent work increasingly addresses dynamic optimization in the presence of transactions costs, which may be substantial in applications to housing (Bayer et al. 2016, Bishop and Murphy 2019, Kennan and Walker 2011).

the no-interference assumption: The outcome (price) of an untreated housing unit in the market is affected by the fact that other housing units were treated with changes to their amenities. In the presence of interference, a policy scenario has an "indirect effect" even on untreated units plus a "direct effect" of treatment on the treated. Consequently, we must make a distinction between the effect of a treatment *scenario* and the effect of treatment *status* for a unit, given the scenario. To analyze such effects, we can draw on extensions to the potential outcome framework to consider effects defined by an entire policy—that is, by a change in *g* anywhere.

In theory, changing the treatment status of just house *h* by itself could itself have general equilibrium effects, which would muddle the distinction between the direct and indirect effects. Assumption A1 rules out this problem:

ASSUMPTION A1 (Local Non-interference). Let $\mathbf{p}_{-h}$ be the vector of prices for all houses except some house h. For all treatment scenarios and $\forall\ h$, $\mathbf{p}_{-h}$ is invariant to whether h is treated.

This assumption, which can be taken as a limiting case of Palmquist's (1992) localized externality, facilitates the distinction between the effect status of treatment on a particular house and the general equilibrium effects of the program.[3]

Let g_h^a be the value of *g* at house *h* which is realized under some potential scenario *a* at *t*=1 and let $\mathbf{g}_{-h}^a$ be the (*H*-1)-dimensional vector of *g* at all houses *except h* in scenario *a*. Let $\mathbf{x}_h^a(\mathrm{g}_h^a)$ be the value of **x** at house *h* in scenario *a*, which itself is a function of g_h^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 $a*$. With this notation, different scenarios *a* describe different possible distributions of *g* at *t*=1.

I incorporate the violation of SUTVA by allowing for different potential prices at house *h* based not only on g_h, but also based on the entire policy vector **g**. Under Assumption A1, we can write the potential outcome for house *h* if we were in the counterfactual state (with no houses

[3] For similar reasons, Hudgens and Halloran (2008) impose a randomization assumption fixing the number of treated units, and VanderWeele and Tchetgen (2011) propose an alternative definition of the direct effect that no longer decomposes the total effect. The local non-interference assumption provides an alternative way to address this issue.

treated) as $p_h^{a'}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'}(\mathrm{g}_h^{a'}))$. The a' superscript on p indicates that the functional relationship depends on the state of the world (in this case the counterfactual state). In Figure 1, this price is illustrated by p_A, from the counterfactual price function evaluated at $\mathrm{g}_h^{a'}$. The potential outcome for house h, if the rest of the market were under policy $a*$, can be written as $p_h^{a^*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a^*}(\mathrm{g}_h^{a'}))$ if house h were not treated (receiving its counterfactual $\mathrm{g}_h^{a'}$) and as $p_h^{a^*}(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a^*}(\mathrm{g}_h^{a*}))$ if house h were treated (receiving g_h^{a*}). These are depicted in Figure 1 by prices p_B and p_C respectively.

As always, causal effects of a policy must compare the policy scenario, a^*, to a counterfactual, a'. Two distinctions here should be emphasized about this comparison which differ from other common settings. First, there may be exogenous changes over time. Thus, in general, the equilibrium under a' is not the same as in t=0. Even if a' is "no program" to change g, the distribution of g may be evolving. Too, even if g would not have changed, p_A is not necessarily equal to p_h^0, as there could be other changes in the economic environment affecting the price function or $\mathbf{x}_h$. Second, with interference, there are endogenous changes from the policy that require untreated units to be evaluated in the policy scenario as well as the counterfactual. Despite the fact that house h is untreated, $p_B = p_h^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a^*}(\mathrm{g}_h^{a'}))$ is not necessarily equal to $p_A = p_h^{a'}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'}(\mathrm{g}_h^{a'}))$ because the treatments at the other houses affect equilibrium prices at all houses, including h.

Following Hudgens and Halloran (2008) and VanderWeele and Tchetgen Tchetgen (2011), define the individual *total effect* (TE) of policy $a*$ for some house h as

$$TE_h(a^*) \;=\; p_h^{a^*}\big(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a^*}(\mathrm{g}_h^{a*})\big) - \; p_h^{a'}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'}(\mathrm{g}_h^{a'})).$$

The total effect is the overall effect of treatment by the policy at house h. In Figure 1, it is equal to p_C - p_A. It can be decomposed into two parts. The individual *indirect effect* (IE) of treatment $a*$ is:

$$\mathrm{IE}_h(a^*) \;=\; p_h^{a^*}\big(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a^*}(\mathrm{g}_h^{a'})\big) - \; p_h^{a'}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'}(\mathrm{g}_h^{a'})).$$

IE_h represents the effect on the price of untreated houses due to the shifting price function between scenarios a' and $a*$. It is the result of interference: Even if house h is untreated, its price may be affected by spillovers from treatments elsewhere. IE_h is depicted in Figure 1 by p_B - p_A.

The individual *direct effect* (DE) of treatment $a*$ for h, conditional on the program going forward in the rest of the market, is defined as

$$\mathrm{DE}_h(a^*) \;=\; p_h^{a^*}\big(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a^*}(\mathrm{g}_h^{a*})\big) - \; p_h^{a^*}\big(\mathrm{g}_h^{a\prime}, \mathbf{x}_h^{a^*}(\mathrm{g}_h^{a\prime})\big).$$

DE_h represents the effect of re-assigning house *h* from an untreated to a treated state, while holding constant the treatment status of the other houses. It is depicted in Figure 1 by p_C - p_B.

As written, TE, IE, and DE all include any effects mediated through changes in **x**. For example, improvements in *g* might motivate households to improve the house in other (observable) ways, or trigger resorting with new households changing the house. Variants of these treatment effects that net out the portions mediated through changes in **x** can be defined for all three. Define the *total unmediated effect* (TUE) and the *direct unmediated effect* (DUE) at $\tilde{\mathbf{x}}_h$ respectively by:

$$\mathrm{TUE}_h(a^*) \;=\; p_h^{a^*}\big(\mathrm{g}_h^{a^*}, \mathbf{x}_h^{a*} = \tilde{\mathbf{x}}_h\big) - \; p_h^{a\prime}(\mathrm{g}_h^{a\prime}, \mathbf{x}_h^{a\prime} = \tilde{\mathbf{x}}_h)$$

$$\mathrm{DUE}_h(a^*) = p_h^{a^*}(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*} = \tilde{\mathbf{x}}_h) - \; p_h^{a^*}(\mathrm{g}_h^{a\prime}, \mathbf{x}_h^{a\prime} = \tilde{\mathbf{x}}_h).$$

TUE and DUE are the same as TE and DE respectively, except they hold $\mathbf{x}_h$ constant. DUE and DE both can be identified by DD hedonic methods. As I show below, DUE is the causal concept with the clearest welfare interpretation as a lower bound on welfare. It represents moving *h* from an untreated to a treated state, while otherwise holding constant the treatment program and holding constant $\mathbf{x}_h$. Accordingly, henceforth I shall focus primarily on DUE.

These individual effects have their respective group averages. Define the average total and direct unmediated effects, respectively, as:

$$\overline{\mathrm{TUE}(a^*)} \;=\; \frac{1}{H}\sum\nolimits_h \mathrm{TUE}_h(a^*),$$

$$\overline{\mathrm{DUE}(a^*)} \;=\; \frac{1}{H}\sum\nolimits_h \mathrm{DUE}_h(a^*).$$

More generally, we could define similar averages over any sub-set of houses $\mathcal{H}'$. In the case of the direct unmediated effect we then write

$$\overline{\mathrm{DUE}_{\mathcal{H}'}(a^*)} \;=\; \frac{1}{\sum_h 1(h \in \mathcal{H}')}\sum\nolimits_{h\in\mathcal{H}'} \mathrm{DUE}_h(a^*).$$

A particular special case is where $\mathcal{H}'$ is simply the subset of treated houses. Then $\overline{\mathrm{DUE}_{\mathcal{H}'}(a^*)}$ is the average direct unmediated effect on the treated, or $\overline{\mathrm{DUET}(a^*)}$.

3. Difference-in-Difference Hedonic Models Identify the Direct Effect

Either $\overline{\text{TE}(a^*)}$ or $\overline{\text{DE}(a^*)}$, or their unmediated variants, could be defined as a "capitalization effect." $\overline{\text{TE}(a^*)}$ captures both the treatment on *h* and the shifting hedonic price function. If we wanted to forecast the effects of the program on prices, relative to a counterfactual of no program, either $\overline{\text{TE}(a^*)}$ or $\overline{\text{TUE}(a^*)}$ would be useful measures. However, the impact on prices *qua* prices are of limited interest for measuring welfare, except for understanding distributional effects on individuals who pay them or receive them as income. As Kuminoff and Pope (2014) and Klaiber and Smith (2013) have emphasized, $\overline{\text{TUE}(a^*)}$ is not the average WTP 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 shift.

Moreover, without additional assumptions $\overline{\text{TE}(a^*)}$ and $\overline{\text{TUE}(a^*)}$ cannot be identified anyway. Because $p_h^{1,a'}$ is not observed, we cannot know the indirect effects.[4] However, $\overline{\text{DUE}(a^*)}$ can still be identified because, by definition, it does not reference scenario a'. It only requires the weaker assumption, typical of DD estimators, that changes from p_h^0 to $p_h^{1,a*}$ for untreated units are the same, on average, as what they would have been for treated units were they left untreated, but the policy went forward. In practice, it allows researchers to use data on observed changes.

The basic argument can be seen in Figure 1, if we replace $p^{a'}(\ )$ with $p^0(\ )$. A treated unit and its matched control both start at p_A. In scenario a^*, the untreated unit moves to p_B, and under the identifying assumption so would the treated unit (in expectation). In this example, at untreated units g does not change, but that assumption is unnecessary.[5] Treated units have an additional shock to g, ending up at p_C. Thus, the identified effect from DD 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, or DUE. As shown

[4] 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, so that, if a' were "no policy," then p_h^0 could be substituted for $p_h^{a'}(g_h^{a'})$ in the expression for TE. This may be especially plausible under regression discontinuity designs. Another assumption is that observations are available at other, untreated markets which can identify the counterfactual time trend. See Hudgens and Halloran (2008) and Manski (2013). Crépon et al. (2013) illustrate the idea.

[5] g can change for the untreated, so long as it would have changed in the same way for the treated absent treatment. E.g., Chay and Greenstone (2005) and Bento et al. (2015) consider price effects of air quality improvements triggered by regulatory thresholds. One can identify the direct effects of the regulations, still allowing for national time trends.

in Section 4, this effect is a lower bound on general equilibrium welfare.

3.1 Identification and Estimation of Capitalization Effects: The Linear Case

For simplicity, consider first a linear model. For any individual house *h*, the *ex-ante* and $a*$ hedonic price functions and their difference are, respectively:

$$p_h^0 = \alpha^0 + \beta^0 g_h^0 + \gamma^{0\prime}\mathbf{x}_h^0 + \xi_h + \varepsilon_h^0, \tag{3a}$$

$$p_h^{a^*} = \alpha^{a^*} + \beta^{a^*} g_h^{a^*} + \gamma^{a^*\prime}\mathbf{x}_h^{a^*} + \xi_h + \varepsilon_h^{a^*}, \tag{3b}$$

$$dp_h^{a*} = d\alpha^{a^*} + d\beta^{a^*} g_h^0 + \beta^{a^*} dg_h^{a^*} + d\gamma^{a^*\prime}\mathbf{x}_h^0 + \gamma^{a^*\prime} d\mathbf{x}_h^{a^*} + d\varepsilon_h^{a^*}, \tag{3c}$$

where in Equation (3c) the differences are taken from the baseline, not from the unobserved counterfactual. This model is a variant of the "generalized difference-in-differences estimator" recommended by Kuminoff et al. (2010). Note the local non-interference assumption A1 is implicitly embedded in (3b), as the parameters are independent of $g_h^{a^*}$.

As written, this model assumes true panel data at the micro level, which in the case of housing is the "repeat sales" model. However, in practice, it could be applied to repeated cross-sections, essentially treating the panel as being at the level of some group (e.g. communities) and using fixed effects in lieu of first differencing.[6] The application in Section 5 illustrates the approach.

In any case, if the hedonic price function does not change between (3a) and (3b), then we can suppress time-scenario superscripts in the parameters and Equation (3c) collapses to

$$dp_h^{a*} = \beta dg_h^{a^*} + \gamma' d\mathbf{x}_h^{a^*} + d\varepsilon_h^{a^*}. \tag{4}$$

In this case, it is clear that DD 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, when the hedonic function does shift, the true model is (3c), so (4) of course is mis-specified. Kuminoff and Pope (2014) refer to this problem as "conflation bias," as the estimates conflate marginal WTP at a point in time (i.e. β^0 or β^{a^*}) with changes in the 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

[6] In that case, we would simply replace equation (3c) with the following fixed effects variant:

$p_h^t = \alpha^0(1\text{-}D^1) + \alpha^{a^*}D^1 + \beta^0 g_h^0(1\text{-}D^1) + \beta^{a^*} g_h^{a^*} D^1 + \gamma^{0\prime}\mathbf{x}_h^0(1\text{-}D^1) + \gamma^{a^*\prime}\mathbf{x}_h^{a^*} D^1 + \xi_c + \varepsilon_h^t,$

where D^1 is an indicator for an observation from the second period. The equation simply stacks equations (3a) and (3b) with period-specific coefficients and a fixed effect for community *c*.

the model, as in Equation (3c), the linear model potentially can identify β^{a^*}, the ex-post marginal WTP 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 may raise additional estimation issues. While allowing for an unobserved time invariant effect, ξ_h, Equations (3) still require a conditional zero mean assumption on $d\varepsilon$ to estimate the full set of parameters. Unfortunately, $d\varepsilon$ may well be correlated with g^0: For example, houses near landfills 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: ($d\varepsilon \perp dg \mid g^0, \mathbf{x}^0, d\mathbf{x}$). This would allow identification of β^{a^*}, even if β^0 were biased.

The fact that it is the *ex post* hedonic price parameter, β^{a^*}, that is identified under the weaker assumptions is the crucial point here. In the context of the linear model, this parameter represents $\overline{\text{DUET}}$ (per unit g), the direct effect netting out the mediated effect of any changes in $\mathbf{x}$. The product $\beta^{a^*}\cdot dg^{a^*}$ is the movement along the ex post hedonic price function in the dimension of g.

3.2 Estimating Direct Effects Without Linearity

The preceding insight extends to nonlinear models like matching estimators as well. To simplify the exposition, consider the case of a binary treatment, which occurs only in the second period: $g^0=0$, $g^{a^*} \in \{0,1\}$.[7] Examples might include cleanup of Superfund sites (Gamper-Rabindran and Timmins 2013, Greenstone and Gallagher 2008), discovery of a cancer cluster (Davis 2004), closing large polluting facilities (Currie et al. 2015, Mastromonaco 2015), arrival of a sex offender (Linden and Rockoff 2008) and so forth. Relaxing the linearity inherent in Equations (3), we can define the potential outcomes using the following semi-parametric assumptions:

(5a) $$p_h^0 = \boldsymbol{\gamma}^{0\prime}\mathbf{x}_h^0 + \varepsilon_h^0,$$

(5b) $$p^{a^*}\left(g_h^{a^*} = 0\right) = \boldsymbol{\gamma}_{g^1=0}^{a^*\prime}\mathbf{x}_h^{a^*} + \varepsilon_{g^{a^*}=0,h}^{a^*},$$

(5c) $$p^{a^*}\left(g_h^{a^*} = 1\right) = \boldsymbol{\gamma}_{g^1=1}^{a^*\prime}\mathbf{x}_h^{a^*} + \varepsilon_{g^{a^*}=1,h}^{a^*},$$

[7] The model of this sub-section could be extended to include multi-valued or even continuous treatments as suggested by Imbens (2000) and Hirano and Imbens (2004). See, e.g., Muehlenbachs d (2015) for a hedonic application.

where the $\boldsymbol{\gamma}$ vectors include an intercept term and $E[\varepsilon^{a^*}_{g^t,h}|\mathbf{x}]$ need not be zero. Equation (5b) represents houses that are not treated ex post, Equation (5c) those that are. This model controls for $\mathbf{x}$ with regressions that differ by treatment status, but allows the effect of g, which is embedded in the ε's, to have any arbitrary form (e.g. Heckman et al. 1997). It again implicitly relies on the local non-interference assumption A1, as $\boldsymbol{\gamma}^{a^*}_{g}$ does not depend on whether any one house h is treated.

This model requires a conditional mean independence assumption on differences in unobservables, slightly weaker than the linear model. For example, if we want to know the effect of the observed policy a^* relative to some counterfactual, then we want to estimate the Average Treatment on the Treated and we require the following assumption.

ASSUMPTION A2 (Conditional mean independence in differences for the treated):

$$E[\varepsilon^{a^*}_{g^1=0} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 1] \;=\; E[\varepsilon^{a^*}_{g^1=0} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 0]. \tag{6}$$

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

Under these conditions, plus the usual requirement of overlapping support, a conditional DD estimator can identify the $\overline{\text{DUE}}$ for the treated. This is stated in the following lemma.

LEMMA. Given A1, A2, and the model of Equations (5),

$$\begin{aligned}
&E\left[\left(p^{a^*}(g^{a^*}=1) - \boldsymbol{\gamma}^{a^*}_{g^1=0}{}'\mathbf{x}^{a^*}\right) - \left(p^0 - \boldsymbol{\gamma}^{a^*}_{g^1=0}{}'\mathbf{x}^0\right) \Big| \, \mathbf{x}^0, g^{a^*}=1)\right] \\
&\quad - E\left[\left(p^{a^*}(g^1=0) - \boldsymbol{\gamma}^{a^*}_{g^1=0}{}'\mathbf{x}^{a^*}\right) - \left(p^0 - \boldsymbol{\gamma}^{a^*}_{g^1=0}{}'\mathbf{x}^0\right) \Big| \, \mathbf{x}^0, g^{a^*}=0)\right] \\
&= E\left[\left(\left(p^{a^*}(g^{a^*}=1) - p^{a^*}(g^1=0)\right) | \mathbf{x}^0, g^{a^*}=1\right)\right. \\
&\quad \left. - \boldsymbol{\gamma}^{a^*}_{g^1=0}{}'\left((d\mathbf{x}|\mathbf{x}^0, g^{a^*}=1) - (d\mathbf{x}|\mathbf{x}^0, g^{a^*}=0)\right)\right] \;=\; \overline{\text{DUET}(a^*)}.
\end{aligned} \tag{7}$$

Proof: The first equality follows immediately from Equations (5) and Assumption A2. The second follows from Equations (5) and Assumption A1. The third follows from the definition of DUET.

$\overline{\text{DUET}(a^*)}$ might be estimated using a regression-adjusted DD matching estimator (e.g. Heckman et al. 1997). Muehlenbachs et al. (2015) and Haninger et al. (2017) use this approach in hedonic applications. For the linear case, the parameter β^1 represents the marginal contribution of *g* along the ex post hedonic, holding constant any effects mediated through **x**. The estimand defined in the Lemma recovers an analogous effect, for those houses treated by the policy, but using weaker econometric assumptions.

3.3. Extensions

The following section is devoted to the economic interpretation of this estimand. Before turning to that discussion, two comments are in order. First, $\overline{\text{DUE}}$ on the treated, which is identified using DD hedonics, arguably identifies effects at the most salient part of the distribution, certainly when evaluating the policy that went into effect. Nevertheless, as shown in appendix A.1, we also can identify treatment effects at untreated units, which would be useful for evaluating counterfactual policies—so long as the observed price function is still representative of the ex post equilibrium. For example, we can always compare the actual policy to an *alternative* counterfactual.

Second, in some cases one may not want to impose the conditional mean independence assumption. Although it is weaker than those required for the standard OLS model, one may be concerned that changes in unobservables are correlated with changes in *g*. If so, one could invoke additional exclusion restrictions and use instrumental variables. For example, Chay and Greenstone (2005) and Bento et al. (2015), considering hedonic regressions of housing prices on air quality, 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, the exogeneity of *g* can be made more credible using regression discontinuity (RD) designs (e.g. Cellini et al. 2010, Chay and Greenstone 2005, Gamper-Rabindran and Timmins 2013, Greenstone and Gallagher 2008, and Lang 2018). In principle, RD designs can be used in either panel or cross-sectional settings. In the neighborhood of the discontinuity, treatment is as good as randomly assigned, as in a controlled experiment (Lee and Lemieux 2010). This has two key implications for interpreting hedonics. First, the local average treatment effect is the same as the average treatment effect on the treated, so RD designs identify either. Second, RD estimates are unbiased even when researchers do not include a full set of controls, so long as they do not jump discretely along with the policy (Bayer et al. 2007), though they are less efficient. This suggests it may not always be necessary to control for changes in **x** not caused by the policy (see Appendix A.4 for further discussion).

4. The Welfare Interpretation of Capitalization Effects

The previous section showed that quasi-experimental hedonic studies can identify $\overline{\text{DUET}(a^*)}$, the direct unmediated effect on the treated, if they allow for changing hedonic functions where appropriate. $\overline{\text{DUET}(a^*)}$ is a well-defined economic concept. It is the difference along the ex post hedonic price function between the values of a house at the new and old levels of the amenity respectively, or equivalently the cumulated marginal values along the price function (not demand function), netting out effects mediated through $d\mathbf{x}$: $p^{a^*}(g^{a^*}, \mathbf{x}^0) - p^{a^*}(g^0, \mathbf{x}^0) = \int_{g0}^{a*} \frac{\partial p^{a*}(g, \mathbf{x}^0)}{\partial g} dg$.

What is the economic interpretation of this estimand? In this section, I show that, summed over houses, it represents a lower bound on the total welfare effects of the policy for *all* households, accounting for non-marginal changes in g and general equilibrium price effects, mobility responses, and endogenous responses to **x**. The specific measure bounded is the Hicksian equivalent surplus (ES) for all these effects. Equivalent surplus is similar to the equivalent variation, but differs insofar as it measures the willingness to accept (WTA) to forego the *realized* change in g, in contrast to the g's that would be chosen when foregoing the price change.[8]

The argument for this lower bound on ES is quite simple in a partial equilibrium context where there are no supply or implicit price effects on **x** and no effects on profits, so that the only effects are the change in the distribution of g (Griffith and Nesheim 2013). By a simple revealed preference argument, the household consuming g^1 could save expenditures amounting to $\int_{g^0}^{g^{a*}} \frac{\partial p^{a*}}{dg} dg$ 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.

[8] In other words, it is the change in money that holds utility constant at ex post levels for a change in g, or equivalently, the area under the Hicksian demand curve for g (evaluated at ex post utility) between g^0 and g^1. Whereas equivalent variation is more appropriate for price changes, ES often is used for exogenous changes in quantities or qualities (see Hicks 1943 or Freeman et al. 2014, Ch. 3).

This can be seen in Figure 2, which depicts a Marshallian bid function for g intersecting the two marginal price functions at the chosen quantities g^0 and g^1. It also depicts a Hicksian demand function evaluated at u^1, which intersects the Marshallian function at g^1. Hicksian ES is the entire shaded area under the Hicksian demand function from g^0 to g^1. (In contrast, equivalent variation would be the area under the function from the level of g chosen under p^0 to achieve u^1, depicted here as $g(p^0(\,), u^1)$, up to g^1.) Clearly, $\int_{g0}^{ga*} \frac{\partial p^{a*}}{dg} dg$, in darker shading, is a lower bound on ES = $\int_{g0}^{ga*} h(g, u^1) dg$. The argument is analogous to the fact that a Paasche index is a lower bound on the value of a quantity change.

In fact, under Assumption A1 (local non-interference) and an additional assumption of non-negative profits, this bound remains true even in general equilibrium with an endogenous change in the entire hedonic price function, endogenous changes in the supply of **x**, and resorting of households. The required assumption on profits is:

ASSUMPTION A3 (non-negative profits). The change in aggregate profits due to adjustments in **x** from their counterfactual level are non-negative when evaluated at the counterfactual level of g:

$$\sum_h \left[\left(p^{a*}(g_h^{a\prime}, \mathbf{x}_h^{a*}) - p^{a*}(g_h^{a\prime}, \mathbf{x}_h^{a\prime}) \right) - \left(c(g_h^{a\prime}, \mathbf{x}_h^{a*}) - c(g_h^{a\prime}, \mathbf{x}_h^{a\prime}) \right) \right] \geq 0.$$

This assumption, which nests a zero-profit condition, seems intuitive: landlords would not make the investment to change **x** unless it increased profits.

Defining H^T as the number of treated houses, we can state the following Proposition, the key result of this section.

PROPOSITION 1. Given A1 and A3, $H^T \cdot \overline{\mathrm{DUET}(a^*)} \leq$ (ES + Δprofits) for an exogenous change in the distribution of g. The result holds even when hedonic prices, households, and landlords adjust to the change endogenously, with these effects included in the welfare measure. It also holds when there are other changes in the economic environment potentially shifting the price function or **x** over time, but does not include these in the welfare measure.

Proof: See Appendix B.

The proposition states that $\overline{\mathrm{DUET}(a^*)}$, times the number of treated units, is a lower bound on welfare for *all* households, whether affected directly or indirectly (by the changing price function). Note that, although panel data are used to control econometrically for unobservables, $\overline{\mathrm{DUET}(a^*)}$ is the average movement along the ex post hedonic function and ES uses the ex post expenditure function. Thus, only the ex post situation is relevant for the interpretation. If demands or tastes change, the result remains valid, but the evaluation is from the perspective of ex post preferences.

The formal proof follows the outline of a verbal argument in Bartik (1988), clarifying a few ambiguous points.[9] Denote the expenditure function for household *i* as $e_i(p(\,), u)$ where $p()$ 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 $p()$. Denote the restricted expenditure function as $\tilde{e}_i(p(\mathrm{g}, \mathbf{x}), \mathrm{g}, \mathbf{x}, u)$; it is the solution to the expenditure minimization problem when the household is constrained to choose the bundle $(\mathrm{g}, \mathbf{x})$. The basic idea is to define the welfare measure as follows:

$$\begin{aligned} dW = & \sum\nolimits_i \left[\tilde{e}_i\left(p^{a'}\left(\mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), \mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u_i^{a*}\right) - e_i(p^{a*}(\,), u_i^{a*})\right] \\ & + \sum\nolimits_h \left[\left(p^{a*}(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a'}\left(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'}\right)\right) - \left(c\left(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*}\right) - c\left(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'}\right)\right)\right]. \end{aligned} \tag{8}$$

where $\left(\mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right)$ represent the characteristics, in the counterfactual scenario a', of a house actually occupied by household *i* in the counterfactual scenario. The first term in square brackets is, by definition, ES for the change in *g*. The second term in brackets is the change in landlord profits. It is the change in rents, from the shift in the hedonic price function, adjustments in **x** and exogenous changes in *g*, minus the change in costs, evaluated at baseline levels of *g*.

As shown in Appendix B, using Assumption A.3, this measure is equivalent to:

$$\begin{aligned} dW = & \sum\nolimits_i \left[\tilde{e}_i\left(p^{a*}\left(\mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), \mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u_i^{a*}\right) - e_i(p^{a*}(\,), u_i^{a*})\right] \\ & + \sum\nolimits_h \left[\left(p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'})\right) - \left(c(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'})\right)\right] \\ & + \sum\nolimits_h \left(p^{a*}(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*})\right). \end{aligned} \tag{9}$$

[9] Bartik's argument actually was that the movement along the *ex ante* price function would be an *upper* bound on welfare. He does not link this measure to what can be identified econometrically. As noted in the previous section, it is the ex post price function that is identified with panel data, which provides a lower bound by similar logic. Additionally, Bartik (1988) does not provide a mathematical poof. Kanemoto (1988) provided another related bounding proof. However, his model is quite different, involving capitalization into land values in long-run equilibrium. Furthermore, he shows pre-policy prices are a lower bound on a rather unusual version of compensating variation for a subset of the economy, whereas my bound uses the identified ex post prices as a bound on conventional ES.

Now, in the first line, 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. Additionally, the second line is non-negative by Assumption A3. Meanwhile, the third line is the observed measure, the sum of price changes along the ex post price function holding $\mathbf{x}$ constant. As the desired welfare measure is the observed value plus a positive number, the observed value is less than the change in welfare.

Thus, there is a clear welfare interpretation of capitalization effects in a DD framework. They can identify a lower bound on ES for changes in g. (For decreases in g, the lower bound means the estimate is "too negative," which is the same as an upper bound on the welfare *loss*, in absolute values.) In general, the gap represented by the bound is unknown. However, we can expect that it better approximates the true value for smaller shocks. Indeed, for truly marginal changes, the movement along the price function is the derivative, which is equal to marginal WTP by the tangency condition (Equation [2]). In Appendix C, I present simulations finding the lower bound welfare estimates for improvements in g range from 75% of ES for large changes to 92% for smaller (but non-marginal) changes.

Moreover, in Appendix A, I show that this welfare interpretation extends to more general settings than those used in the main body of the paper. First, in A2, I show that, though for expositional simplicity it was easier to develop the model as a closed city, all the results extend to open cities. Second, in A3, I show that it also extends to a world with transactions costs, as long as those costs are super-additive in a network of multiple moves. Finally, in A4, I consider a more negative result: the unmediated effects considered so far require controlling for changes in $\mathbf{x}$ if they are either correlated with g or caused by the policy. However, when these conditions are violated, additional bounding results apply.

5. Application to Changes in Toxic Air Emissions

In this section, I illustrate DD hedonic studies with an application to changes in exposure to plants emitting toxic air pollutants, a question also considered by Currie et al. (2015) and Mastromonaco (2015). In particular, I estimate a variant of Equation (3c), using two cross sections of individual houses and treating local geographic areas as the panel unit (see note 6). My strategy resembles that of Currie et al. (2015) in spirit. They treat plants as observations, looking at the effect of plant openings and closings on average property values within a 1-mile ring of the plant, relative to the effect at 1-2 miles. In contrast, I have microdata on housing transactions, so treat houses as observations, looking at the effect of a changing *number* of plants within a 2-mile ring of the house, while controlling for conditions at 1-mile grid cells using fixed effects. I also consider controlling for changing conditions using 2-mile-cell-by-year fixed effects. Additionally, whereas Currie et al. assume constant hedonic coefficients (as in Equation 5), to avoid conflation bias I allow the hedonic coefficients to evolve between the two time periods, as in Equation (3c). Thus, this approach identifies $\overline{\mathrm{DUE}(a^{*})}$ and the lower bound on welfare.

The specific application is to the Los Angeles area (all of LA and Orange Counties and portions of Riverside, San Bernardino, and Ventura Counties), between 1995 and 2000. Data on toxic air emissions come from the US EPA's TRI database. As discussed by Currie et al. (2015) and Mastromonaco (2015), TRI data are good at identifying polluting plants, but exhibit measurement error in emission levels. Accordingly, I focus on the extensive margin of whether a plant is emitting at all, rather than emission levels. These comings and goings of plants too can be measured with error, if plants fall below the TRI reporting threshold rather than actually shut down. Currie et al. overcome this problem by merging TRI data to confidential data on plants' operations. Unfortunately, those data are not available to me. However, I approximate their approach by coding a plant as operating at year *t* if it appears in the TRI database at *any point* between *t* and *t*-8 and *t* and *t*+8. Thus, plants that come, go, and return from the TRI reports are assumed to be operating throughout the period. (Eight years takes the 1995 data back to the beginning of the TRI program.)

Exposure to TRI facilities was imputed to a house in two ways. My main approach weights facilities by their distance, with a facility having a weight of max{0, 1- ½*Distance*}, where *Distance* is the distance in miles from the facility to a given census block. Thus, e.g., a facility 2 or more miles away from a census block receives a weight of zero, a facility 1 mile away has a weight of ½, and a facility co-located in the block has full weight. As an alternative, I use the raw count of facilities within two miles. The top panel of Table 1 gives summary statistics for these measures, by year. It shows a decline in exposure to TRI, with the average house experiencing an 8% reduction in the number of facilities. Finally I take the square root of these counts plus 1. This transformation approximates more complicated polynomials while restricting the function to be monotonic in pollution.

Data on real estate transactions were acquired from Fidelity National Data Service, a market research firm providing proprietary data. They include the sales price, date of the sale, number of bedrooms, number of bathrooms, square footage, lot size, year built, and the census block in which it is located. After restricting the data to single family homes and arm's length transactions, and after discarding certain outliers, the data include nearly 140,000 observations.[10] The bottom panel of Table 1 summarizes the housing data. Housing values increased over the period and the housing stock aged, but other characteristics remained fairly constant.

Typically, researchers use census tracts or zip codes as geographic units for constructing spatial fixed effects. However, these geographic units may evolve over time. Additionally, they are based on population density, so their sizes vary widely. This is problematic if small geographies are systematically more homogeneous than large ones, so that there is more variation with which to estimate effects from the latter, and if large geographies also vary systematically in unobserved ways from other areas. Finally, geographies creating homogeneous areas (like Census tracts) may inflate the spatial scale of very localized effects by systematically aggregating the affected area to similar areas nearby. For all these reasons, arbitrary zones like a 1-mile grid are preferable to census geographies when controlling for spatial fixed effects (see Banzhaf and Walsh 2008). Accordingly, I impose a 1 square mile grid over the area, using the grid cells for spatial fixed effects.

Identifying $\overline{\mathrm{DUE}(a^*)}$ requires changes in the number of active TRI sites to be orthogonal to changes in unobserved factors affecting prices, after conditioning on baseline conditions and fixed

---

[10] I drop all observations with housing prices below \$50,000 or above \$2m, with lot sizes smaller than 1000 sq. feet or larger than 10 acres, with living area smaller than 500 sq ft or larger than 5,000 sq ft, with zero bathrooms or more than 7.5 bathrooms, or more than 10 bedrooms. At the low end, these observations likely reflect either coding errors or non-primary residences; at the high end, they represent extremely grand houses, where mis-specification of a linear regression is likely to pose problems. These dropped observations account for about 2.7% of the data.

effects. To gauge the plausibility of this assumption, I look at pre-existing price trends. In particular, I regress 1990-1995 prices on *future* 1995-2000 changes in exposure interacted with a time trend, plus contemporaneous exposure and the hedonic variables listed in Table 1. Table 2 shows the results for the coefficient of interest, the interaction between the time trend and the future change in facilities. It shows the change, in percentage points, in annual housing appreciation over 1990-95, for each additional plant to which the house would become exposed between 1995 and 2000. The coefficients are negative and statistically significant at conventional levels for the unweighted model, and marginally so for the weighted model. However, they are not economically meaningful. A one-plant increase in future exposure reduces 1990-95 appreciation by 0.0001 to 0.0003 percentage points per year, or less than $1 at the mean of the data. I also consider more flexible variants of this model, interacting cubic time trends with each of four categories of changes in future exposure: Increases, no change, a 1-plant decrease, and a decrease of 2+ plants. Figure 3 shows the results. As seen in the graph, housing prices in LA were declining in the early 1990s after a long boom, but note that the price levels are not monotonic by category: Prices in areas experiencing no change in pollution are higher than those showing either increases or decreases. More importantly, the trends are fairly parallel—remarkably so for the last three years. This evidence reassuringly suggests that TRI sites were not closing in areas that were already gentrifying.

For the main model of interest, I regress log price on the square root of nearby TRI sites, cubic functions of lot area, living area, bathrooms, bedrooms, and age, plus time invariant fixed effects for the grid cell, and year-quarter dummies to pick up proportionate housing price inflation. Table 3 presents the results using distance-weighted exposure. The table shows results for six separate regressions in separate columns, which can be grouped into three pairs. The three pairs differ by the set of fixed effects: 1-mile grid fixed effects, 1-mile grid fixed effects plus county-year effects, or 2-mile grid-by-year fixed effects. Within each pair, Model A represents the conventional approach with a time-invariant hedonic price function (i.e., Equation(4)), whereas Model B uses the time-varying approach advocated in this paper.

The first three rows present the pollution coefficients from the hedonic regressions, with standard errors, clustered by grid cell, in parentheses. All pollution coefficients are negative and statistically significant. Nevertheless, they should be interpreted with care. As discussed above, the coefficient in the time-invariant model likely suffers from conflation bias. Additionally, in the time-varying model the 1995 coefficients are only identified under the assumption that changes in unobservables are uncorrelated with baseline conditions, and anyway are not used to construct the lower bound on ES. The 2000 coefficients are identified under the weaker assumption that changes in unobservables and changes in exposure are independent, conditional on baseline exposure and other variables. Interestingly, for all three models, the time-invariant coefficient falls within the interval bracketed by the corresponding two time-varying coefficients, although as discussed by Kuminoff and Pope (2014) this pattern need not always hold. The 2000 coefficient is higher than the time-invariant model, so in this application it appears conflation bias is downward.

The next panel shows the respective capitalization effects from the 1995-2000 change in TRI exposure, using the 2000 price function in the case of Models B. The estimates shown, in $ billions, aggregate the predicted price increases, re-weighting the data so as to reflect the number of owner-occupied housing units in each census tract, to obtain the total estimated capitalization for the LA owner-occupied housing stock. When using Model B and the ex-post function to measure the effects, this measure represents $\overline{\text{DUET}}$ for the observed 1995-2000 change in exposure, which is the lower bound on the general equilibrium ES. The standard errors (parentheses) and confidence intervals (brackets) are based on a cluster bootstrap. The bootstrap accounts for sampling error in estimation, forecasting (i.e. predicting direct effects), and reweighting tracts.

Using Models A, the conventional approach yields capitalization effects ranging from $5.8b to $6.5b. Using Models B, $\overline{\text{DUET}}$ ranges from $6.7b to $7.5b. These estimates represent a present value for what presumably is a permanent shock to amenities. With about 5.4m owner-occupied housing units in the study area, and at a discount rate of 5%, these values for $\overline{\text{DUET}}$ work out to $62-$69 per LA household per year. Thus, in this case, even the lower bound measure is substantial and may be informative for policy. The last row of the panel directly compares the estimates of $\overline{\text{DUET}}$ from columns B to the conventional capitalization approach in columns A. For this application, the conventional approach gives estimates that are \$0.8b to \$1.2b, or 14-20 percent, higher. Moreover, the differences are statistically significant. Thus, not only does the approach suggested in this paper bring clarity to the estimand as a lower bound on welfare, in this case it also is attended by a tighter bound. Finally, comparing across models, we see that allowing for time-varying fixed effects (as in Models 2 and 3) does not appear to close the gap or reduce conflation bias: Time-varying coefficients are required.

The last panel of Table 3 offers another welfare measure often presented, the linearized marginal WTP. The approach computes the derivative of the hedonic function at each point, interpreted as marginal WTP, and multiplies this marginal value by the change in g. Typically, these marginal values are evaluated at the baseline points, so for Models B they now are off the 1995 function. (In the case of Models A, these estimates would mechanically be the same as those from the previous panel if the hedonic function were linear; any differences are due to non-linearities.) We now have the reverse result: the lower coefficients for 1995 in the time invariant model result in lower estimates of marginal WTP. However, these estimates cannot be directly compared to the capitalization results, for three reasons. First, whereas (for Models B) the capitalization estimates are a lower bound on a general equilibrium version of ES, linearized marginal WTP is an upper bound on a partial equilibrium version of Hicksian compensating surplus (CS). As ES is likely higher than CS, the "upper bound" can still be lower. Second, they use the 1995 coefficients instead of the 2000 coefficients. Third, those coefficients are only identified under stronger assumptions.

These general patterns are insensitive to alternative modelling strategies, including using the raw count of TRI sites instead of the distance-weighted count and using 0.5 or 2-mile grid cells instead of 1-mile grid cells. These results are available upon request.

6. Conclusions

For decades, economists have used the hedonic model to estimate demands for the implicit characteristics of differentiated commodities, including otherwise unpriced local public goods and amenities. The traditional cross-sectional approach to hedonic estimation has recovered marginal WTP for amenities when unobservables are conditionally independent of the amenities, but has been criticized as biased when this condition is not met (Greenstone 2017).

In response, economists have introduced panel econometric models using DD and related approaches to identify capitalization effects. Unfortunately, the interpretation of these effects has not been clearly perceived in the literature, perhaps because there is a range of meanings to the word "capitalization" and "causal effect" when the price function shifts. In this paper, I show that DD and related hedonic methods can 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 further show that this is a lower bound measure on Hicksian ES. Simulations suggest the lower bound provides valuable information, on the order of 75% to 92% of the actual equivalent surplus.

These results have two implications for hedonic research. First, researchers can employ DD and other quasi-experimental methods to overcome estimation problems, while maintaining a clear connection to the underlying economic model, a point that sometimes has been challenged. Second, from those estimates they can recover bounds on non-marginal welfare measures, without turning to Rosen's (1974) "second stage" or more structural approaches.

Future work might consider how quasi-experimental methods might be extended to account for price and distributional effects. For example, Crépon et al. (2013), Hudgens and Halloran (2008), and Manski (2013) propose ways to identify indirect effects using variation in treatment programs across groups (markets, in the hedonic context), while still identifying direct effects from variation in treatment assignment within a group. Such methods might allow researchers to identify total price effects, and hence transfers among subpopulations of buyers and sellers. As Sieg et al. (2004) discuss, such price changes can have important distributional welfare effects.

Additional work might consider ways to average different bounds to improve the approximation. Some quasi-experimental strategies can plausibly identify effects in multiple cross-sections, especially when they use RD or IV strategies. Examples include Greenstone and Gallagher (2008), Gamper-Rabindran and Timmins (2013), Gopalakrishnan et al. (2011), and Haninger et al. (2017). When they do, there is at least the potential to identify separate effects in two or more cross sections. If the movement along the hedonic price function can be estimated in the ex ante period as well as the ex post, it would be possible to construct an *upper* bound analogous to the lower bound discussed here (Bartik 1988). If so, it may be further possible to average these effects to get a second-order approximation to welfare, as suggested by Banzhaf (2020).

References

Abadie, A. 2005. Semiparametric Difference-in-Differences Estimators. *Review of Economic Studies* 72(1): 1-19.

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. 2020. Panel Data Hedonics: Rosen's First Stage as a "Sufficient Statistic." *International Economic Review* 61: 973-1000.

Banzhaf, H.S., and R.P. Walsh. 2008. Do People Vote with their Feet? An Empirical Test of Tiebout. *American Economic Review* 98(3): 843-63.

Bartik, T.J. 1988. Measuring the Benefits of Amenity Improvements in Hedonic Price Models. *Land Economics* 64(2): 172-83.

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. 2016. A Dynamic Model of Demand for Houses and Neighborhoods. *Econometrica* 84(3): 893-942.

Bento A., M. Freedman, and C. Lang. 2015. Who Benefits from Environmental Regulation? Evidence from the Clean Air Act Amendments. *Review of Economics and Statistics* 97(3): 610-22.

Bishop, K.C., and A.D. Murphy. 2019. Valuing Time-Varying Attributes Using the Hedonic Model: When is a Dynamic Approach Necessary? *Review of Economics and Statistics* 101(1): 134-45.

Bishop, K.C., and C. Timmins. 2019. Estimating the Marginal Willingness to Pay Function without Instrumental Variables. *Journal of Urban Economics* 109: 66-83.

Cellini, S.R., F. Ferreira, and J. Rothstein. 2010. The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design. *Quarterly Journal of Economics* 125(1): 215-61.

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

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.

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

Epple, Dennis, Louis Quintero, and Holger Sieg. 2020. A New Approach to Estimating Hedonic Equilibrium Models for Metropolitan Housing Markets. *Journal of Political Economy* 128(3): 948-83.

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. 2017. The Continuing Impact of Sherwin Rosen's Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy* 125(6): 1891-1902.

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.

Haninger, K., L. Ma, and C. Timmins. 2017. The Value of Brownfield Remediation. *Journal of the Association of Environmental and Resource Economists* 4(1): 197-241.

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., R.L. Matzkin, and L. Nesheim. 2010. Nonparametric Estimation of Nonadditive Hedonic Models. *Econometrica* 78(5): 1569-91.

Freeman, A.M. III, J.A. Herriges, and C.L. Kling. 2014. *The Measurement of Environmental and Resource Values*. RFF Press.

Hicks, J.R. 1943. The Four Consumer Surpluses. *Review of Economic Studies* 11(1): 31-41.

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: 981-9.

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., 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. 2014. Do "Capitalization Effects" for Public Goods Reveal the Public's Willingness to Pay? *International Economic Review* 55(4): 1227-50.

Lang, C. 2018. Assessing the Efficiency of Local Open Space Provision. *Journal of Public Economics* 158: 12-24.

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.

Mastromonaco, R. 2015. Do Environmental Right-to-Know Laws Affect Markets? Capitalization of Information in the Toxic Release Inventory. *Journal of Environmental Economics and Management* 71: 54-70.

Muehlenbachs, L., E. Spiller, and C. Timmins. 2015. The Housing Market Impacts of Shale Gas Development. *American Economic Review* 105(12): 3633-3659.

Palmquist, R.B. 1992. Valuing Localized Externalities. *Journal of Urban Economics* 31(1):59-68.

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.

Pope, J.C. 2008. Fear of Crime and Housing Prices: Household Reactions to Sex Offender Registries." *Journal of Urban Economics* 64(3): 601-14.

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

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.

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.

Figure 1. Defining Capitalization Effects

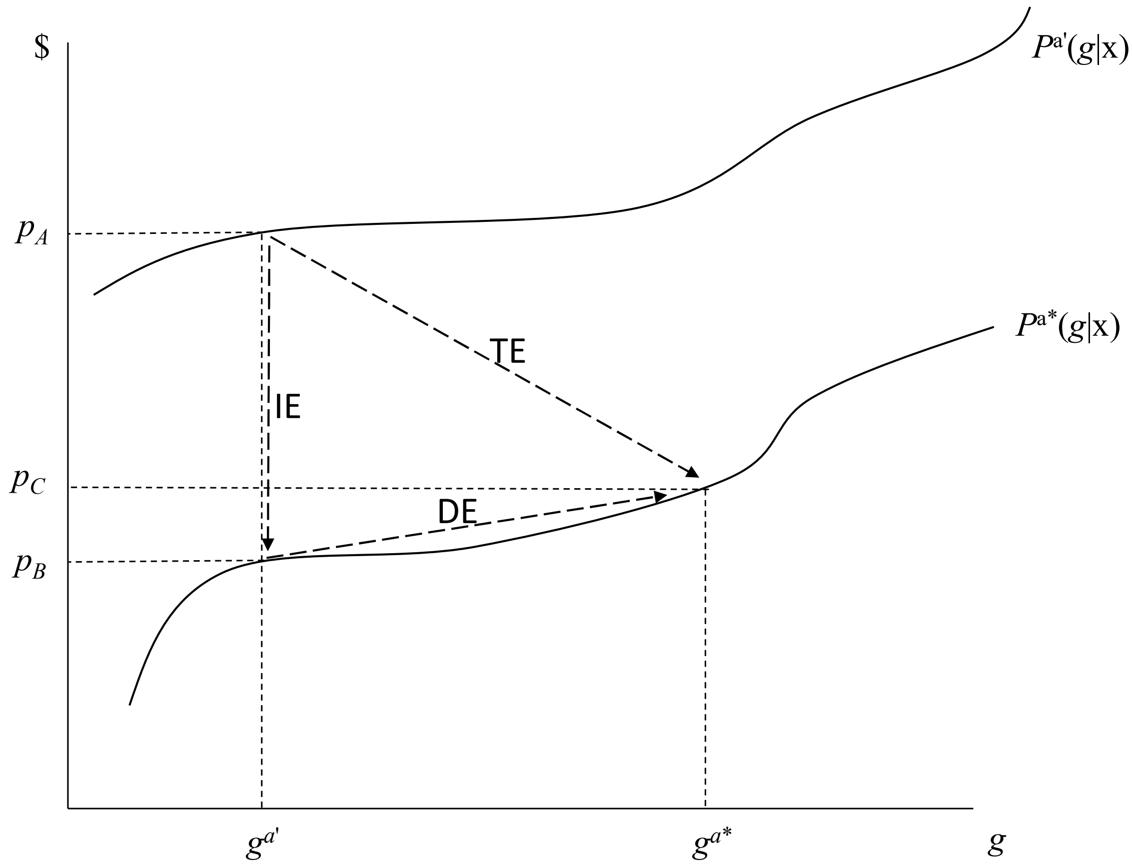

Figure 1 illustrates the indirect effect (IE), direct effect (DE), and total effect (TE) of a policy shifting the distribution of an amenity g. The policy shifts the hedonic price function from $p^{a'}(\,)$ to $p^{a*}(\,)$. Even the price of untreated houses are affected by this shift, moving from p_A to p_B, which is IE. Treated houses move from p_A to p_C, which is TE for these units. This total effect can be decomposed into IE+DE.

Figure 2. Bounds for ES for Changes in Characteristics

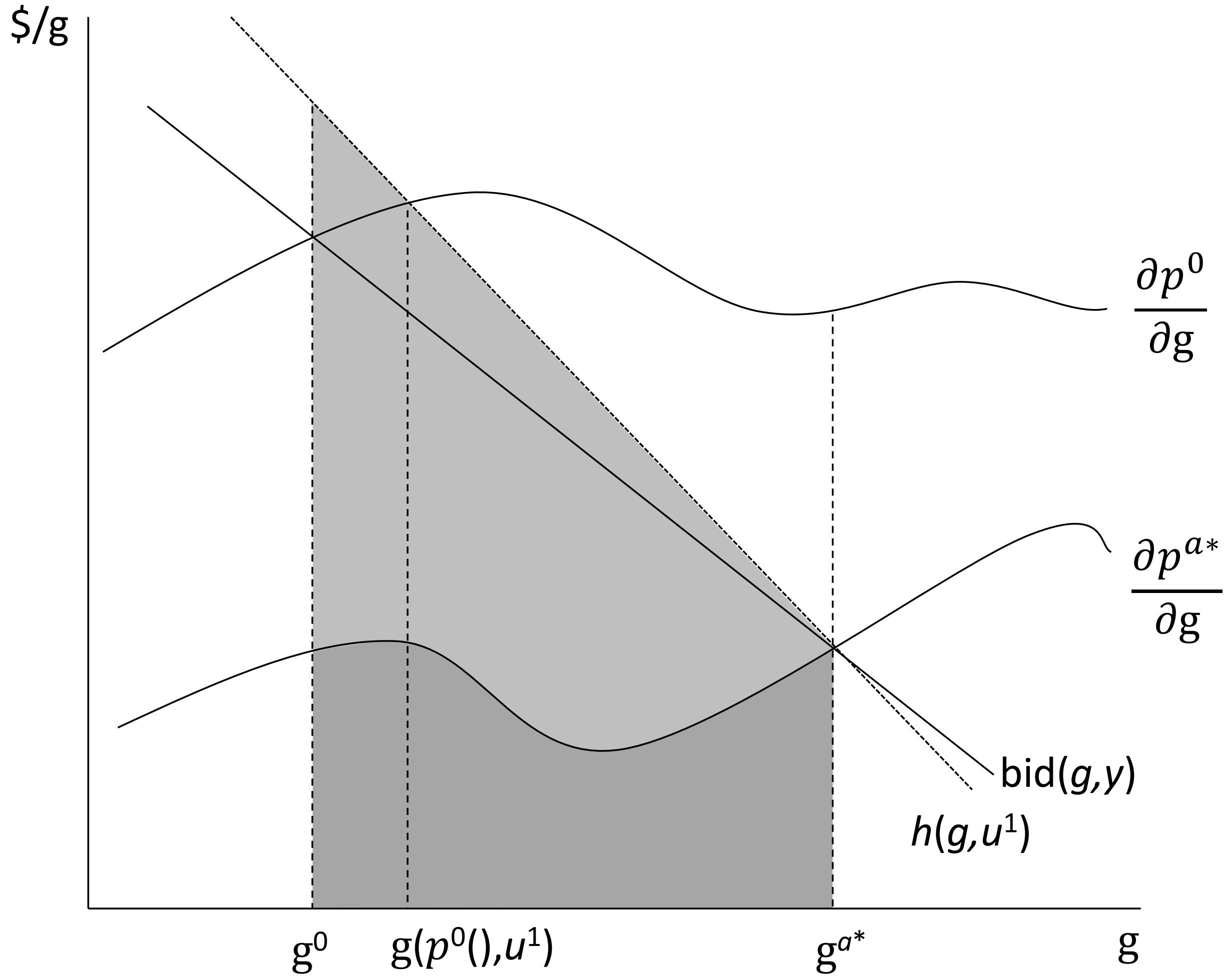

Figure 2 shows the Hicksian Equivalent Surplus (ES) as the area under the Hicksian demand curve $h(\ )$ from g^0 to g^{a*}. A partial-equilibrium version of the lower bound is illustrated by the fact that this area exceeds the movement along the price function, or the area $\int_{g^0}^{g^{a*}} \frac{\partial p^{a*}}{dg} dg$. The text shows this bound extends to general equilibrium.

Figure 3. Pre-existing Trends in Housing Prices.

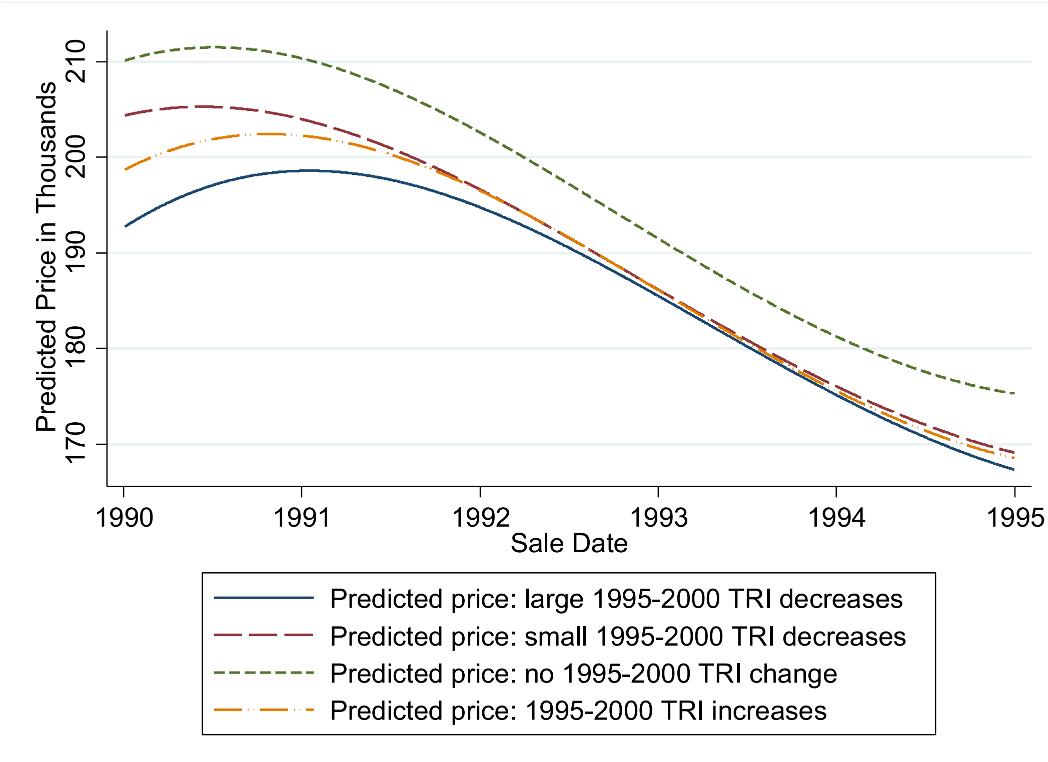

Figure 3 shows 1990-1995 mean predicted prices for four categories of houses, those with 1995-2000 decreases in TRI exposure of 2+ plants, a decrease of one plant, no change, and increases in expsoure.

Table 1. Summary Statistics of Housing Amenities, by Year

	1995 Mean (Std. Dev.)	2000 Mean (Std. Dev.)
A. Measures of TRI Exposure		
No. TRI facilities within 2 miles	2.50 (4.31)	2.29 (3.80)
Distance-weighted No. TRI facilities within 2 miles	0.74 (1.44)	0.68 (1.3)
B. Structural Characteristics		
Sale Price	193,736 (129,691)	270,914 (190,782)
Lot Size (sq. feet)	9,617 (14,285)	9,979 (16,736)
Living Area (sq. feet)	1,672 (640)	1,723 (701)
Bedrooms	2.05 (0.73)	2.09 (0.82)
Bathrooms	3.20 (0.84)	3.22 (0.92)
Age	31.13 (20.86)	36.15 (21.71)

Table 2. Pre-Existing Trends in Housing Prices

	Weighted Number of Facilities within 1 mile (1)	Number of Facilities within 1 mile (2)
Time Trend x 1995-2000 Change in facilities	-0.000272 (0.000150)	-0.000123 (0.0000465)
R^2	0.78	0.78

This table shows the coefficients from regressing housing transactions between 1990 and 1995 on 1995-2000 changes in the number of TRI facilities within 1 mile interacted with a time trend. The regressions control for time trends (without interaction), contemporaneous TRI exposure, the hedonic variables listed in Table 1, and 1-mile grid-cell fixed effects. Standard errors clustered by grid cell reported in parentheses.

Table 3. Results of Application to TRI data.

	Model 1 1 mi grid FE		Model 2 1 mi grid FE + County-Year FE		Model 3 2 mi grid–Year FE	
	Model 1A	Model 1B	Model 2A	Model 2B	Model 3A	Model 3B
Estimated Coefficients						
Sqrt TRI facilities (Time Invariant)	-0.1313 (0.0137)		-0.1375 (0.0137)		-0.1466 (0.0111)	
Sqrt TRI facilities (1995)		-0.1047 (0.0213)		-0.0962 (0.0221)		-0.0526 (0.0103)
Sqrt TRI facilities (2000)		-0.1493 (0.0156)		-0.1653 (0.0161)		-0.1657 (0.0135)
Capitalization Effects						
Movement along price function, $b [95% CI]	5.917 (0.569) [4.802 - 7.033]	6.732 (0.622) [5.513 - 7.952]	6.200 (0.601) [5.023 - 7.378]	7.440 (0.710) [6.048 - 8.831]	6.378 (0.478) [5.442 - 7.314]	7.495 (0.579) [6.360 - 8.630]
Difference from Model A [95% CI]		0.815 (0.367) [0.095 - 1.535]		1.239 (0.423) [0.411 - 2.068]		1.117 (0.170) [0.784 - 1.450]
Linearized Marginal WTP						
Marginal Value x d*g*, $b [95% CI]	5.833 (0.583) [4.690 - 6.977]	4.650 (0.823) [3.036 - 6.264]	6.108 (0.589) [4.954 - 7.261]	4.274 (0.840) [2.628 - 5.921]	6.510 (0.525) [5.480 - 7.539]	2.334 (0.394) [1.563 - 3.106]
Difference from Model A [95% CI]		-1.183 (0.513) [-2.188 - -0.177]		-1.834 (0.628) [-3.065 - -0.602]		-4.175 (0.567) [-5.286 - -3.064]
R^2	0.751	0.753	0.780	0.782	0.816	0.818

Each column represents a separate regression. Each pair of columns (Models1-3) uses different fixed effects. For each model, Model A imposes a time-invariant hedonic price function on TRI facilities and all hedonic controls, while Model B allows all coefficients to vary over time. The 2nd panel shows the movement along the hedonic price functions, using the ex post function for the year-specific models, which gives the lower bound. The 3rd panel shows the product of marginal value times the observed 1995-2000 change in TRI exposures. Standard errors in parentheses are clustered by grid cells. Confidence intervals for welfare measures, in square brackets, are cluster bootstrapped.

Appendix A. Extensions and Discussion

A.1. Effects at Other Parts in the Distribution

Estimating the average direct effect on the treated makes most sense for an ex-post welfare evaluation of a policy. However, in principle, one could imagine asking other questions of the data. For example, perhaps one would want to know the direct effect of treatment for some other subset of houses $\mathcal{H}'$ under an alternative policy. That is, given the equilibrium a^*, we might want to know the average effect on prices of changing treatment status for that subset. For such questions, one would merely adapt Assumption A2 to establish the conditional independence assumption needed to estimate $\overline{\text{DUE}_{\mathcal{H}'}(a^*)}$, a mix of the usual treatment-on-the-treated and treatment-on-the-controls estimands.[11]

Under the appropriate assumptions, then, one could claim to identify the movement along the observed ex post hedonic function $p^{a^*}(\,)$ for any house in the set $\mathcal{H}'$ if its g were to have changed by some specified level. Thus, the results of Section 3 generalize easily. However, the effects identified would still be the movement along $p^{a^*}(\,)$.

This movement along $p^{a^*}(\,)$, even for some margin of the distribution $\mathcal{H}'$ that differs from the actual ex post treatments, can be interpreted in two ways. One possibility is that we are considering an alternative ex post distribution of g, evaluated relative to the same baseline a'. For small tweaks from the actually observed ex post distribution, under which the observed price function is plausibly the same, we could continue to use the results of Section 4 to interpret the results as a lower bound for the welfare effects of such an alternative policy. However, for a very different policy, treating a very different set $\mathcal{H}'$, we would expect a different ex post hedonic price function. While a movement along that alternative price function would still be a lower bound on ES, without

---

[11] Specifically, the assumptions required would be:

$$E[\varepsilon^{a^*}_{g^1=0} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 1, h \in \mathcal{H}'] = E[\varepsilon^{a^*}_{g^1=0} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 0]$$

$$E[\varepsilon^{a^*}_{g^1=1} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 1] = E[\varepsilon^{a^*}_{g^1=1} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 0, h \in \mathcal{H}'].$$

If $\mathcal{H}'$ is the entire sample then these assumptions collapse to the usual assumptions needed for an Average Treatment Effect (in this case, average direct unmediated treatment effect).

observing it such a lower bound cannot be constructed. This is simply a recognition of the fact that ex post policy evaluation requires ex post data.

But another way the set of treated houses $\mathcal{H}'$ may differ is to imagine an alterative counter-factual equilibrium a' to which we are comparing the actual ex post equilibrium a^*. For example, one might one to compare the welfare effects of the actual policy to a smaller policy, in which a subset $\mathcal{H}''$ would have been treated anyway. The true ES would be different, then, but the ex post hedonic would remain unchanged and the lower bound can still be constructed from it. One would simply evaluate DUE for the right set of houses. In this way, the lower bound results of Section 4 also extend to other margins of the distribution.

A.2. Open Cities.

For expositional reasons it has been useful to think in terms of a closed region with a constant set of houses and households. As briefly noted earlier, however, that assumption plays no role in the results of the paper. Whether there are other cities (unaffected by the policy directly, that is, with no exogenous change in their distribution of g), with households endogenously moving in and out, would, of course, affect the true welfare measure ES. It also would affect the distributional impacts of the policy. Importantly, though, even if prices in the other cities are affected by the policy in the study region, there is no exogenous change in g in those cities and, presumably, no endogenous changes. In the language of Section 3, with open cities there may be indirect effects elsewhere, but no direct effects. And it is only those direct effects that are used to construct the lower bound.

A.3. Transactions Costs.

In many hedonic applications, such as the regular, repeated purchase of computers, we might plausibly take the choice to purchase a new model as exogenous to the change in attributes. In other cases, such as the housing example emphasized in this paper, changes in available attributes might cause people to re-optimize to a new $\{g, \mathbf{x}\}$ bundle. In the housing setting as well as other contexts, such as automobile purchases, the transactions costs of changing attribute bundles is not trivial.

The analysis of Section 4 can be extended to include such mobility or other transactions costs. In the case when there are no transactions costs, we can compute hypothetical compensations at alternative prices and a given utility level, without specifying the cost-minimizing solution at the actual price giving rise to that utility level. Transactions costs, however, create path dependency. Thus, we must at least specify a starting point from which the consumer moves when re-optimizing, because transactions costs will be a function of the initial as well as the final point. In particular, in Equation (8), when evaluating the hypothetical compensation $\tilde{e}_i$ equivalent to the policy shock (when the consumer is constrained to be at the counterfactual bundle), we take the perspective of moving the consumer from their optimal ex post bundle. That is, when we compute the hypothetical expenditure required to maintain ex post utility at the counterfactual price level and the counterfactual bundle $\left(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right)$, the required expenditure *includes any money needed to pay transactions costs when moving back to the bundle consumed in a'*. In the case of an improvement, the notion here is that the consumer's willingness to accept compensation to forego the improvement requires paying transaction costs to move away from the improved point and back to a particular bundle $\left(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right)$.

Denote the utility achieved in the ex post scenario $a*$, when households have to pay transactions costs moving from some reference point *R*, as $u_i^{TC(R),a*}$. Let $e_i^{TC(R)}(p,u)$ be the minimal cost of achieving some level of utility in the presence of transactions costs, when those transaction costs are incurred from a specific reference point *R*. Then, in the first line in Equation (8), we replace the expression

$$\sum_i \left[\tilde{e}_i\left(p^{a'}\left(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}\right) - e_i\left(p^{a*}(\,), u_i^{a*}\right)\right] \equiv ES$$

with the expression

$$\sum_i \left[\tilde{e}_i^{TC(a*)}\left(p^{a'}\left(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{TC(R),a*}\right) - e_i^{TC(R)}\left(p^{a*}(\,), u_i^{TC(R),a*}\right)\right] \equiv ES^{TC}.$$

The term $e_i^{TC(R)}\left(p^{a*}(\,), u_i^{TC(R),a*}\right)$ is the expenditure required to achieve the ex post utility with transactions costs, $u_i^{TC(R),a*}$, when the consumer actually does have to pay those transactions costs. The term $\tilde{e}_i^{TC(a*)}\left(p^{a'}\left(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{TC(R),a*}\right)$ is, as before, the expenditure required to maintain the ex post utility when constrained to be at the counterfactual bundle and given counterfactual prices, but with two differences. First, now transactions costs must be paid from the ex post point a^* and figured into the required expenditure level. Second, the utility level at which it is evaluated also differs.

As shown in the proof in Apendix B, all other terms also remain unchanged, so the argument flows through without any other alterations. We are left in the end with the expression:

$$(10)\quad \begin{aligned} dW = & \sum_i \left[\tilde{e}_i^{TC(a*)}\left(p^{a*}\left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), \mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{TC(R),a*}\right) - e_i^{TC(R)}\left(p^{a*}(\,), u_i^{TC(R),a*}\right)\right] \\ & + \sum_h \left[\left(p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'})\right) - \left(c(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'})\right)\right] \\ & + \sum_h \left(p^{a*}(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*})\right). \end{aligned}$$

If it were not for the transactions costs paid, this expression would be the same as Equation (9) evaluated at a different utility level. Consequently, most of the argument flows through. However, we must account for the fact that the reference points for the two expenditure functions in Equation (10) differ: One involves $TC(a*)$ whereas the other involves $TC(R)$. In Appendix B, I show the lower bound result flows through with the following additional assumption.

ASSUMPTION A4 (super-additivity). Let $TC(a, a')$ be the transactions cost of moving from a bundle a to a bundle a'. Then $TC(a, a'') \leq TC(a, a') + TC(a', a'')$.

This mild assumption states that the transactions cost of moving directly from bundle a to a'' is no higher than the transactions cost of moving first from a to a' and then from a' to a''.

With this additional assumption, we have the following proposition.

PROPOSITION 2. Given A1, A3, and A4, $H^T \cdot \overline{\mathrm{DUE}_{\mathcal{H}'}(a^*)} \leq (\mathrm{ES}^{\mathrm{TC}} + \Delta\text{profits})$ for an exogenous change in the distribution of g for a set of houses $\mathcal{H}'$. The result holds even when hedonic prices, households, and landlords adjust to the change endogenously, with these effects included in the welfare measure. It also holds when there are other changes in the economic environment potentially shifting the price function or the levels of $\mathbf{x}$ over time, but it does not include these in the welfare measure.

Proof: See Appendix B

Thus, the main results of this paper apply even in the case of transaction costs.

A.4. The Direct Mediated Effect.

The analysis in the body of the paper focused on the direct unmediated effect, in which we control for changes in **x**. But some characteristics may be unobserved and some recent hedonics papers have intentionally omitted contemporaneous characteristics (while controlling for baseline levels), presumably to include mediated effects. What if we do not control for such changes? The consequences differ for two cases. First, if the changes in unobservables are simply correlated with the change in *g* globally, but not actually caused by it, then any quasi-experimental design where the source of variation in d*g* is exogenous may overcome this problem, such as an RD or other IV design.

However, changes in **x** may well be caused by the policy-induced changes in *g,* either through resorting (Bayer et al. 2007) or through changes in the housing stock induced by the changes in *g*. Unfortunately, in that case the observed measure would incorporate the gross benefits of general equilibrium adjustments to **x**, without netting out the costs of providing them. This can undermine the lower bound interpretation, although we can still establish sufficient conditions under which it holds. To see this, rewrite Equation (9) as:

$$\begin{aligned} dW = & \sum_i \left[\tilde{e}_i\left(p^{a*}\left(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}\right) - e_i(p^{a*}(\,), u_i^{a*})\right] \\ & - \sum_h \left(c(g_h^{a'}, \mathbf{x}_h^{a*}) - c\left(g_h^{a'}, \mathbf{x}_h^{a'}\right)\right) \\ & + \sum_h \left(p^{a*}(g_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(g_h^{a'}, \mathbf{x}_h^{a'})\right). \end{aligned} \tag{11}$$

Here, the first set of terms remains the same as in Equation (9), as do the change in costs, while the last set of terms is the direct effect gross of any change in **x**, which is now what we observe. For the lower bound to still hold (i.e. for the last set of terms to be less than *dW*), we would need

$$\begin{aligned} & \sum_h \left(c(g_h^{a'}, \mathbf{x}_h^{a*}) - c(g_h^{a'}, \mathbf{x}_h^{a'})\right) \leq \\ & \qquad \sum_i \left[\tilde{e}_i\left(p^{a*}\left(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}\right) - e_i(p^{a*}(\,), u_i^{a*})\right]. \end{aligned} \tag{12}$$

Essentially, the change in costs from the change in **x,** which are associated with the change in *g*, cannot be too big; more precisely, it cannot be bigger than the ES for households from the changes in *g*. A sufficient, but not necessary, condition is that they are non-positive. One plausible example is where the unobserved **x** is literally free. For example, *g* might be transportation infrastructure, which could affect unobserved changed in air quality.

Nevertheless, this analysis suggests under the most general case it is important to control for changes in **x** whenever possible to have the cleanest welfare interpretation. It is worth emphasizing here that the question centers on changes in **x**, not levels. Unobserved, unchanging levels of **x** cancel out in the comparison, which motivate the use of DD strategies at the outset.

APPENDIX B. PROOF OF PROPOSITIONS.

Proof of Proposition 1.

As noted in the text, our measure of the change in welfare is:

$$dW = \sum_i \left[\tilde{e}_i\left(p^{a'}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right) - e_i(p^{a*}(\,), u^{a*}_i)\right]$$
$$\text{(8)} \qquad + \sum_h \left[\left(p^{a*}(g^{a*}_h, \mathbf{x}^{a*}_h) - p^{a'}\left(g^{a'}_h, \mathbf{x}^{a'}_h\right)\right) - \left(c\left(g^{a'}_h, \mathbf{x}^{a*}_h\right) - c\left(g^{a'}_h, \mathbf{x}^{a'}_h\right)\right)\right].$$

The right side of Equation (8) can be decomposed as follows:

$$dW = \sum_i \left[\tilde{e}_i\left(p^{a*}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right) - e_i(p^{a*}(\,), u^{a*}_i)\right]$$
$$\text{(13)} \qquad + \sum_i \left[\tilde{e}_i\left(p^{a'}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right) - \tilde{e}_i\left(p^{a*}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right)\right]$$
$$+ \sum_h \left(p^{a*}(g^{a*}_h, \mathbf{x}^{a*}_h) - p^{a*}\left(g^{a'}_h, \mathbf{x}^{a*}_h\right)\right) + \sum_h \left(p^{a*}\left(g^{a'}_h, \mathbf{x}^{a*}_h\right) - p^{a*}\left(g^{a'}_h, \mathbf{x}^{a'}_h\right)\right)$$
$$+ \sum_h \left(p^{a*}(g^{a'}_h, \mathbf{x}^{a'}_h) - p^{a'}(g^{a'}_h, \mathbf{x}^{a'}_h)\right) - \sum_h \left(c(g^{a'}_h, \mathbf{x}^{a*}_h) - c(g^{a'}_h, \mathbf{x}^{a'}_h)\right).$$

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

$$dW = \sum_i \left[\tilde{e}_i\left(p^{a*}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right) - e_i(p^{a*}(\,), u^{a*}_i)\right]$$
$$\text{(14)} \qquad + \sum_h \left[\left(p^{a*}(g^{a'}_h, \mathbf{x}^{a*}_h) - p^{a*}(g^{a'}_h, \mathbf{x}^{a'}_h)\right) - \left(c(g^{a'}_h, \mathbf{x}^{a*}_h) - c(g^{a'}_h, \mathbf{x}^{a'}_h)\right)\right]$$
$$+ \sum_h \left(p^{a*}(g^{a*}_h, \mathbf{x}^{a*}_h) - p^{a*}(g^{a'}_h, \mathbf{x}^{a*}_h)\right)$$
$$+ \sum_i \left[\left(\tilde{e}_i\left(p^{a'}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right) - \tilde{e}_i\left(p^{a*}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right)\right)\right.$$
$$\left. - \left(p^{a'}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right) - p^{a*}\left(g^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right)\right)\right].$$

Note that, 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 expression simplifies to:

$$dW = \sum_i \left[\tilde{e}_i\left(p^{a*}\left(\mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), \mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{a*}_i\right) - e_i(p^{a*}(\,), u^{a*}_i)\right]$$

(9)
$$+ \sum_h \left[\left(p^{a*}(\mathrm{g}^{a'}_h, \mathbf{x}^{a*}_h) - p^{a*}(\mathrm{g}^{a'}_h, \mathbf{x}^{a'}_h)\right) - \left(c(\mathrm{g}^{a'}_h, \mathbf{x}^{a*}_h) - c(\mathrm{g}^{a'}_h, \mathbf{x}^{a'}_h)\right)\right]$$

$$+ \sum_h \left(p^{a*}(\mathrm{g}^{a*}_h, \mathbf{x}^{a*}_h) - p^{a*}(\mathrm{g}^{a'}_h, \mathbf{x}^{a*}_h)\right).$$

But, in the first line, 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. Additionally, the second line also is non-negative by Assumption A3. Thus,

(15)
$$\sum_h \left(p^{a*}(\mathrm{g}^{a*}_h, \mathbf{x}^{a*}_h) - p^{a*}(\mathrm{g}^{a'}_h, \mathbf{x}^{a*}_h)\right) \leq dW.$$

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

Note that Assumptions A1 and A3 are necessary but not sufficient conditions for the bound, in the sense that the proposition is not an if-and-only-if statement.

Proof of Proposition 2

Our measure of the change in welfare is now:

$$dW = \sum_i \left[\tilde{e}^{TC(a*)}_i \left(p^{a'}\left(\mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}\right), \mathrm{g}^{a'}_{i(a')}, \mathbf{x}^{a'}_{i(a')}, u^{TC(R),a*}_i\right)\right.$$

(16)
$$\left. - e^{TC(R)}_i\left(p^{a*}(\,), u^{TC(R),a*}_i\right)\right]$$

$$+ \sum_h \left[\left(p^{a*}(\mathrm{g}^{a*}_h, \mathbf{x}^{a*}_h) - p^{a'}(\mathrm{g}^{a'}_h, \mathbf{x}^{a'}_h)\right) - \left(c(\mathrm{g}^{a'}_h, \mathbf{x}^{a*}_h) - c(\mathrm{g}^{a'}_h, \mathbf{x}^{a'}_h)\right)\right].$$

It is similar to Equation (8), but now $u^{TC(R),a*}_i$ is the utility achieved in scenario $a*$ when the household has to pay transaction costs from the reference pointy *R*. Additionally, the expenditure function $e^{TC(R)}_i\left(p^{a*}(\,), u^{TC(R),a*}_i\right)$ is the same as in Equation (8) but it takes into account the expenditure necessary to pay the transaction cost. Likewise the constrained expenditure function $\tilde{e}^{TC(a*)}_i(\,)$ takes such expenditures into account.

Repeating the same steps as the proof for Proposition 1, we can derive the following expression:

$$
\begin{aligned}
dW = & \sum_i \left[\tilde{e}_i^{TC(a*)}\left(p^{a*}\left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), \mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, u_i^{TC(R),a*}\right) - e_i^{TC(R)}\left(p^{a*}(\,), u_i^{TC(R),a*}\right)\right] \\
& + \sum_h \left[\left(p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'})\right) - \left(c(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a'})\right)\right] \\
& + \sum_h \left(p^{a*}(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(\mathrm{g}_h^{a'}, \mathbf{x}_h^{a*})\right).
\end{aligned}
\tag{17}
$$

which parallels Equation (9).

To prove the proposition, we must show that the first term, in square brackets, is non-negative for all i, as in the case without transaction costs.

We will make use of two facts. The first is that

$$
\begin{aligned}
u_i^{TC(R),a*} & = u\left((\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*}), y - p^{a*}(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*}) - TC(R, a^*)\right) \\
& \geq u\left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}, y - p^{a*}\left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right) - TC(R, a')\right).
\end{aligned}
\tag{18}
$$

The first line simply states that the utility achieved is the utility from the services of the bundle consumed and expenditure on other goods, which is income minus the cost of the hedonic bundle minus the transaction costs to obtain it when the reference point is R. The second line follows by revealed preference: because $(\mathrm{g}_h^{a*}, \mathbf{x}_h^{a*})$ was chosen to maximize utility given the price function and the transaction costs from reference point R, it must yield higher utility than the bundle $\left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right)$ given the same price function and same reference point for transactions.

The second fact we will use is that

$$
\begin{aligned}
& \tilde{e}_i^{TC(a*)}\left(p^{a*}\left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), \left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), u_i^{TC(R),a*}\right) \\
& \qquad - \tilde{e}_i^{TC(R)}\left(p^{a*}\left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), \left(\mathrm{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'}\right), u_i^{TC(R),a*}\right) \\
& = \left(TC(R, a^*) + TC(a^*, a')\right) - TC(R, a').
\end{aligned}
\tag{19}
$$

This expression merely states that the expenditure needed to achieve a given level of utility at a given bundle of hedonic attributes, under a given price function, but under two different transactions costs, differs only by the level of those transactions costs. In one case, the household is going directly from

R to a', in another indirectly via $a*$.

From these two facts, we can complete the proof through the following steps:

$$\tilde{e}_i^{TC(a*)}\left(p^{a*}\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right),\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right),u_i^{TC(R),a*}\right)$$

$$=\tilde{e}_i^{TC(R)}\left(p^{a*}\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right),\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right),u_i^{TC(R),a*}\right)+\left(TC(R,a^*)+TC(a^*,a')\right)-TC(R,a')$$

$$\geq \tilde{e}_i^{TC(R)}\left(p^{a*}\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right),\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right),u\left(\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right),y-p^{a*}\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right)-TC(R,a')\right)\right)$$
$$+\left[\left(TC(R,a^*)+TC(a^*,a')\right)-TC(R,a')\right]$$

$$=e_i^{TC(R)}\left(p^{a*}(\,),u_i^{TC(R),a*}\right)+\left[\left(TC(R,a^*)+TC(a^*,a')\right)-TC(R,a')\right]$$

$$\geq e_i^{TC(R)}\left(p^{a*}(\,),u_i^{TC(R),a*}\right).$$

The first equality follows from re-arranging Equation (16). Next, the inequality follows from Expression (18), plus the fact that the expenditure function is increasing in u. The next equality follows from the fact that $\tilde{e}_i^{TC(R)}(\,) = e_i^{TC(R)}\left(p^{a*}(\,),u_i^{TC(R),a*}\right)=y$. That is, the expenditure necessary to achieve $u_i^{TC(R),a*}$ when actually paying the prices of scenario $a*$ and the transactions cost from R is just y; likewise, the expenditure necessary to achieve the utility of $\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right)$ and numeraire consumption $y-p^{a*}\left(\mathrm{g}_{i(a')}^{a'},\mathbf{x}_{i(a')}^{a'}\right)-TC(R,a')$, when constrained to be at that hedonic bundle and to pay those prices and transactions costs is just y. The last inequality follows by Assumption A4: the transaction costs of moving directly from R to a' is no higher than that from moving indirectly via $a*$. This completes the proof.

Appendix C. Simulations

Because the true welfare measure can never be known in an actual application, the lower bound discussed in this paper can only be illustrated and assessed using simulations. Accordingly, this section illustrates the paper's findings by simulating hedonic housing equilibria and shocking the equilibria with changes to *g*. Prices are determined in an equilibrium model, the "true" welfare measures calculated, and the lower bound welfare measures estimated using the econometric methods discussed above. Of course, the bounds depend on the underlying parameters assumed.

I simulated equilibria for 100 economies, each with 1000 households and houses. Households have utility over numeraire consumption, a scalar-valued index of local public goods *g*, and a scalar-valued index of housing attributes *x*. Housing bundles are obtained by purchasing a lot of a fixed size with attributes *g* at the hedonic market price, and purchasing housing capital *x* at a constant price *p*. Household's *i*'s preferences for a house at lot *j* can be represented by an indirect utility function of the following form:

(20) $$v_i(p, \mathrm{g}_j, y_i) = \frac{y_i + \gamma \mathrm{g}_j}{\exp\left(\frac{\exp(\alpha_i + \delta_i \mathrm{g}_j) p^{1+\beta}}{1+\beta}\right)}$$

These preferences represent a variant of the repackaging model, a standard approach to modeling differentiated products (see von Haefen 2007). With these preferences, housing demand, conditional on *g*, takes the form:

(21) $$x_i^* | g_j = \exp(\alpha_i + \delta_i \mathrm{g}_j + \beta \ln p + \ln y).$$

Thus, the characteristic of interest, *g*, enters preferences in two ways. First, it enters through the demand for housing, *x*, as a weak complement. Increases in *g*, ceteris paribus, increase the demand for *x* at a location (and the consumer surplus). Second, through the additive term in the numerator of Equation (20), it enters as a perfect substitute to numeraire consumption.

The utility parameters are set as follows: α is uniformly distributed on (-2.2, -1.8), $\beta = -1.1$, δ is triangularly distributed on (0.025, 0.035, 0.07) to allow outliers in tastes for public goods, and

γ=100. Income is log-normally distributed with mean 11.1 and standard deviation 0.4 (and truncated at \$30,000 and \$200,000). These parameters were calibrated so that equilibrium rents (including land and capital expenditures) as a percent of income range from 20% of income at the 10th percentile of the distribution to 41% at the 90th percentile, with a mean of 29%, which approximates the US expenditure shares for housing.

The public good g is uniformly distributed on (1, 3) in the baseline scenario. In the ex post scenarios, either 10%, 25%, or 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 $g^1 = g^0 + (3-g^0)/3 + 1$. Households respond to these policy shocks by re-optimizing. In the base model, they do so without transactions costs, but in alternative models transactions costs are introduced at \$2000, \$10,000, or \$20,000 (converted to a flow using a 5% discount rate).

Finally, equilibrium prices in each scenario were perturbed by a two-component error. One component, ε_{jt}, is normally and independently distributed and calibrated such that the standard deviation of the error is equal to 1% 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 (so it enters the price function orthogonally to g and x). The second component, ξ_j, is fixed over time, but is correlated both with g^1 and with the treatment conditional on g^1. That is, it represents a time invariant factor correlated with public goods and with their improvement. The existence of such a term motivates, econometrically, using DD models.

Panel A of Table C1 shows the results from the main DD estimates (Equation 3c), when there are no transactions costs. The first column shows the true ES across the three scenarios, with "small" the case where 10% of houses are treated (i.e., 100 out of 1,000), "medium" the case with 25%, and "big" 50%. The second column shows a "best estimate" of the lower bound, in which the true fixed effect and true value of housing capital are subtracted from the property value, and the remaining value from the second period fitted to ex post g using local linear regression. These values are not meant to represent a realistic econometric model, but rather to represent the lower bound construct to be estimated. The next four columns give the estimates from four DD models. The models all condition on baseline levels as well as fixed effects, as in the generalized DD estimator suggested by Kuminoff et al. (2010). The first variant is a simple linear model such as Equation (3c), regressing changes in price on changes in *g*, changes in *x*, and baseline *g* and *x*. The second uses a translog model regressing logged price on a quadratic of logged *g* and logged *x*, using both periods, and with fixed effects. The third uses local linear regression in *g*, partialing out the effects of changes in *x* and baseline *g* and *x* (Yatchew 1998). The final model uses an inverse-probability weighted regression model with the double-robust property, first predicting the probability of treatment using a logit model based on a quadratic of baseline *g*, then predicting the change in price from treatment while regression-adjusting for changes in *x*, using the inverse of the first-stage probabilities as weights.

The best estimate of the lower bound is about 92% of the ES value in the small scenario, gradually decreasing as the scenario gets bigger to 75% in the biggest scenario. This pattern makes sense: for small changes, the marginal value approximates all the welfare information, so small changes along the price function can capture it well, whereas for larger changes the approximations fails to capture the curvature in the WTP function as well as any shifts in the price function. Still, the estimates are the right order of magnitude and provide reasonably useful information. The empirical estimates are all relatively close to the best estimate of the lower bound, suggesting they can recover that lower bound and are useful if properly interpreted. The results also are consistent with the simulations of Klaiber and Smith (2013), who find in their own simulation exercise that DD hedonics yield welfare estimates lower than a true measure of welfare (in their case, the sum of compensating variations).

Panels B and C of Table C1 repeat the ES and "best estimate" columns, but report the results under alternative econometric designs. Panel B uses the same basic econometric models as Panel A, but uses only a cross-sectional model with ex-post data. Thus, the ξ_j component of the error is no longer differenced out. Because it is positively correlated with *g*, this omission creates an upward bias in the lower bound measures, which overstate the importance of g in the cross section. This, of course, is the motivation for using difference-in-differences in the first place. (Note the matching model is omitted from this panel, as selection into treatment is based on baseline data, which by assumption are not available.) The table shows that the cells in the final four columns are higher than their counterparts in Panel A, but closer to the true welfare values. However, there is something of a "two wrongs make a right" flavour to the results, as in this case an upper bias on a lower bound brings us closer to the true ES without overshooting it.

Panel C returns to the DD model but now omits the regression adjustment for changes in $\mathbf{x}$. These models represent the direct mediated effect discussed in Appendix A.4. Because g and x are complements, these models, like those of Panel B, increase the estimates of the lower bound. Again, all cells in Panel C are higher than their respective counterparts in Panel A, and all are higher than the "best estimate" of the lower bound. They are no longer lower bounds relative to the true ES, suggesting that Equation (12) is not satisfied here.

Table C2 reports the results when imposing transactions costs, in the case of the medium-sized policy. Again, the first column shows the true ES across the three scenarios, and the second column shows the "best estimate" of the lower bound. Panel A shows the results from the main DD estimates. The first row of the panel repeats the corresponding row from Table C1, with no transactions costs, for comparison. For the remaining rows, the true ES changes to reflect Equation (10) rather than (9), as discussed in Section A.3. First, the ex post utility level is lower, as households have to pay transactions costs if they want to re-optimize; second, the compensation envisioned when placing somebody at their ex ante level of g includes the transactions costs to put them there. In these simulations, the latter effect slightly dominates, with the ES increasing from $1,527 to $1,933 as transactions costs increase. Nevertheless, the patterns are similar to Table C1. The best estimate of the movement along the ex-post hedonic continues to be a lower bound for ES, as do most of the empirical estimates, with the exception of the local linear model, which performs poorly in the high transactions-cost simulations). Panel B shows the results from the cross-sectional models. Again, ignoring the fixed effect biases the hedonic price function, leading to an upward bias on a measure of a lower bound. Ignoring changes in housing capital has a similar effect as in Panel C of Table C1 as well, again with an upward bias on the lower bound measure (these estimates are not reported but are available upon request).

Additional References

von Haefen, R.H. 2007. Empirical Strategies for Incorporating Weak Complementarity into Consumer Demand Models. *Journal of Environmental Economics and Management* 54: 15-31.

Yatchew, A. 1998. Nonparametric Regression Techniques in Economics. *Journal of Economic Literature* 36: 669-721.

Table C1. Results of Simulations

Scenario	Median Value Across 100 Simulations (5^{th} and 95^{th} Percentiles in Parentheses)					
	ES	"Best Estimate" of Lower Bound	Linear	Translog	Semi-parametric Local linear	Semi-parametric Matching
A. Base Differences-in-Differences Model						
Small	\$508 (398 - 588)	\$471 (373 - 537)	\$435 (352 – 498)	\$326 (240 – 424)	\$680 (421 – 1,307)	\$472 (370 – 537)
Medium	\$1,527 (1,370 - 1,685)	\$1,293 (1,173 - 1,393)	\$1,131 (1,006 – 1,221)	\$1,021 (858 – 1143)	\$1,516 (1,157 – 2,157)	\$1,231 (1,119 – 1,331)
Large	\$3,219 (3,054 - 3,406)	\$2,409 (2,303 - 2,508)	\$2,102 (2,019 – 2,223)	\$2,051 (1,872 – 2,274)	\$2,416 (2,090 – 3,064)	\$2,019 (1,926 – 2,115)
B. Second Period Cross-Sectional Model						
Small	\$508 (398 - 588)	\$471 (373 - 537)	\$507 (397 – 607)	\$590 (423 – 763)	\$607 (483 – 730)	N/A
Medium	\$1,527 (1,370 - 1,685)	\$1,293 (1,173 - 1,393)	\$1,435 (1,306 – 1,580)	\$1,658 (1,434 – 1,978)	\$1,686 (1,531 – 1,832)	N/A
Large	\$3,219 (3,054 - 3,406)	\$2,409 (2,303 - 2,508)	\$2,975 (2,822 – 3,167)	\$3,380 (2,904 – 3,891)	\$3,314 (3,101 – 3,479)	N/A
C. Direct Mediated Effects (*x* omitted)						
Small	\$508 (398 - 588)	\$471 (373 - 537)	\$1,276 (1,031 – 1,484)	\$1,160 (950 – 1,386)	\$732 (465 – 1,353)	\$1,298 (1,053 – 1,511)
Medium	\$1,527 (1,370 - 1,685)	\$1,293 (1,173 - 1,393)	\$3,225 (2,912 – 3,493)	\$2,986 (2,662 – 3,263)	\$1,554 (1,2013 – 2,216)	\$3,253 (2,942 – 3,523)
Large	\$3,219 (3,054 - 3,406)	\$2,409 (2,303 - 2,508)	\$6,061 (5,629 – 6,478)	\$5,728 (5,342 – 6,074)	\$2,471 (2,126 – 3,106)	\$5,975 (5,627 – 6,366)

This table shows the welfare effects and bounds in the simulated equilibria, for policies with small, medium, and large changes in *g*. Each cell shows the median value across 100 simulations, plus the 5^{th} and 95^{th} percentiles. The first column shows the true equivalent surplus. The second shows the lower bound as the movement along the (known) price function. The remaining four columns show empirical counterparts to this bound, using econometric estimators. Panel A uses DD methods; Panel B uses only the ex-post cross section (ignoring time-invariant unobservables); and Panel C ignores changes in **x**.

Table C2. Results of Simulations with Transactions Costs

Transactions Cost	Median Value Across 100 Simulations (5^{th} and 95^{th} Percentiles in Parentheses)					
	ES	"Best Estimate" of Lower Bound	Linear	Translog	Semi-parametric Local linear	Semi-parametric Matching
A. Base Differences-in-Differences Model						
None	$1,527 (1,370 - 1,685)	$1,293 (1,173 - 1,393)	$1,131 (1,006 – 1,221)	$1,021 (858 – 1143)	$1,516 (1,157 – 2,157)	$1,231 (1,119 – 1,331)
$2,000	$1,616 (1,469 - 1,776)	$1,220 (1,079 – 1,343)	$891 (742 – 1,014)	$999 (845 – 1,172)	$1,069 (874 – 1,353)	$1,122 (976 – 1,243)
$10,000	$1,803 (1,578 - 2,009)	$866 (512 – 1,117)	$566 (270 – 790)	$710 ($328 – $959)	$2,887 (1,785 – 4,900)	$668 (325 – 938)
$20,000	$1,933 (1,435 - 2,144)	$636 (7 – 934)	$531 (-103 – 797)	$467 (-363 – 765)	$5,206 ($2,904 – 10,073)	$521 (-156 – 885)
B. Second Period Cross-Sectional Model						
None	$1,527 (1,370 - 1,685)	$1,293 (1,173 - 1,393)	$1,435 (1,306 – 1,580)	$1,658 (1,434 – 1,978)	$1,613 (1,433 – 1,829)	N/A
$2,000	$1,616 (1,469 - 1,776)	$1,220 (1,079 – 1,343)	$1,190 (984 – 1354)	$1,705 (1,267 – 2,213)	$1,613 (1,432 – 1,830)	N/A
$10,000	$1,803 (1,578 - 2,009)	$866 (512 – 1,117)	$859 (567 – 1,084)	$1,335 (1,032 – 1,715)	$1,249 (917 – 1,500)	N/A
$20,000	$1,933 (1,435 - 2,144)	$636 (7 – 934)	$834 (213 – 1,085)	$1,194 (598 – 1,411)	$1,028 (393 – 1,316)	N/A

This table shows the welfare effects and bounds in the simulated equilibria, for the policy giving the "medium" change in *g*, under four different transactions costs of zero, $2,000, $10,000, and $20,000. Each cell shows the median value across 100 simulations, plus the 5^{th} and 95^{th} percentiles. The first column shows the true equivalent surplus. The second shows the lower bound as the movement along the (known) price function. The remaining four columns show empirical counterparts to this bound, using econometric estimators. Panel A uses DD methods; Panel B uses only the ex-post cross section (ignoring time-invariant unobservables).