

Difference-in-Differences Hedonics

H. Spencer Banzhaf

Georgia State University, NBER, PERC

August 2019

Abstract

Traditional cross-sectional estimates of hedonic price functions theoretically can recover marginal willingness to pay for characteristics, but face endogeneity problems when some characteristics are unobserved. To help overcome such problems, economists have introduced difference-in-differences and other quasi-experimental econometric methods into the hedonic model. Unfortunately, the welfare interpretation of the estimands has not been clear. This paper shows that, when they condition on baseline data, they identify the "average direct unmediated effect" on prices from a change in characteristics. It further shows that this effect is a lower bound on welfare, specifically Hicksian equivalent surplus plus the change in profits. The paper illustrates these results with an application to toxic facilities' effects on housing prices.

I thank David Albouy, Olivier Beaumais, Kelly Bishop, Jared Carbone, Michael Greenstone, Kerry Smith, Chris Timmins, Reed Walker and Jeffrey Zabel for helpful comments and suggestions on previous drafts. I am especially grateful to Paul Ferraro and Nick Kuminoff for very detailed comments and discussions. I also thank numerous seminar and conference participants. Finally, I thank Megan Lawson and Randy Walsh for sharing some of the data used in Section 7 with me.

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, to control for quality change, and to price underlying characteristics. The standard hedonic model begins with a single cross-section and uses the insight that a household choosing a particular product has a marginal willingness to pay for an attribute equal to the derivative of the hedonic price function with respect to the attribute. While working with cross-sections has been the standard hedonic project for decades, more 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, so that changes to the hedonic attributes of interest are plausibly exogenous.¹ Parmeter and Pope (2013) provide an introduction to and review of this literature, and Kuminoff et al. (2010) illustrate with simulations the importance of using DD hedonic techniques to control for unobservables.

From a standpoint of statistical estimation, this work is clearly an important improvement, allowing for identification of the hedonic price function under much weaker assumptions about unobservable characteristics. However, from an economic perspective, it seemingly has come at the cost of a clear interpretation of the estimand: what economic question it answers is not always clear, or at least has not been perceived clearly in the literature. The ambiguity arises because the hedonic equilibrium is fundamentally cross sectional: Households face a tradeoff among products at a point in time, not across time.

Although the results of this paper apply to any hedonic context, for concreteness, consider the specific application of housing price functions. Let the hedonic price function be $p^t = p^t(g^t, \mathbf{x}^t)$, where p^t is the price (or cost) of the hedonic bundle in period t , g is a continuous characteristic of interest, and \mathbf{x} is a vector of other characteristics. For example, in applications to housing, g might be a public good of interest and \mathbf{x} a vector of housing characteristics (lot size, dwelling size, and

¹ Examples include Bento et al. (2015), Cellini et al. (2010), Chay and Greenstone (2005), Currie et al. (2015), Davis (2004, 2011), Figlio and Lucas (2004), Gopalakrishnan et al. (2011), Greenstone and Gallagher (2008), Haninger et al. (2017), Lang (2018), Linden and Rockoff (2008), Mastro Monaco (2015), Muehlenbachs et al. 2015, and Pope (2008).

so forth). At a point in time, the derivative of a price function with respect to the characteristic of interest, $\partial p'(g^t, \mathbf{x}^t)/\partial g$, is the marginal willingness to pay. However, it is not clear how this marginal willingness to pay can be identified in the context of DD hedonic regressions. The dependent variable in a DD hedonic regression is

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

which mixes information from two equilibria. Recently, the literature has begun to refer to such differences as "capitalization" (Chay and Greenstone 2005, Klaiber and Smith 2013, Kuminoff and Pope 2014, Mastromonaco 2015, Muehlenbachs et al. 2015, Parmeter and Pope 2013), in contrast to the slope of a single equilibrium hedonic price function.

The link between such capitalization and the underlying economic model is not immediately clear. For changes to the set of characteristics for a small subset of houses, the equilibrium hedonic price function can be taken as constant over a short time period, so that DD models can be interpreted within a single equilibrium (Palmquist 1992). But in the more general case, a large change in the supply of an amenity will shift the hedonic price function. Too, other changes in the economic environment (income, other amenities) taking place over the longer time periods used in many studies, such as ten years, would also shift the price function. When the hedonic price function shifts for either reason, panel-data studies compare prices at two different equilibria in potentially confusing ways. The confusion is compounded by ambiguity about the meaning of language borrowed from the program evaluation literature, such as "*the* capitalization effect," in the sense of a causal effect, of a change in amenities on prices, when there are various such effects with differing interpretations.

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

Willig condition (Palmquist 2005 and Smith and Banzhaf 2004).

On the other hand, Klaiber and Smith (2013), Kuminoff and Pope (2014), and Parmeter and Pope (2013) have argued that, because it combines two equilibria, the capitalization effect often answers an ill-defined economic question. Taken together, their arguments essentially involve two points in the presence of changes in the hedonic price function. First, the total price effect of a change in g is not the same as willingness to pay, nor indeed related to it in any clear way. Instead, it conflates willingness to pay (defined within the context of one equilibrium price function) with changes in the equilibrium price function. Second, estimation of $p'(g^t, \mathbf{x}^t)$ will be biased if the general equilibrium effects on the price function are ignored. Both points are correct. Nevertheless, DD hedonic studies can provide meaningful welfare measures that account not only for marginal willingness to pay for g , but also for general equilibrium effects.

This paper clarifies the issue by introducing distinctions made in the treatment effect literature when the Stable Unit Treatment Value Assumption (SUTVA) is violated. Distinguishing between the indirect and direct effects of treatment, it shows that, when properly conditioning on baseline conditions, DD hedonic studies identify the movement along the ex post hedonic price function, which is not the same as the total price effect if the function shifts endogenously. Second, it shows this effect can be interpreted as a lower bound on Hicksian equivalent surplus (ES) (plus the change in profits) for an improvement in g , even in the presence of general equilibrium shifts in the price function and endogenous adjustments to the supply of \mathbf{x} . 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, panel data on household choices are not required, and even repeated cross sections of houses can be used. Thus, as long as changes in the attribute of interest meet a basic conditional independence assumption, a wide variety of data sets can be used to estimate the lower bound, even with repeated cross sections separated over time.

The paper illustrates these results with an application to the value of reduced toxic emissions in southern California between 1995 and 2000. Although the estimates are a lower bound, the estimated value is substantial, at about \$8 billion for the present value of the realized decrease in emissions, or about \$74 per year.

DD hedonics can thus provide useful information about the underlying structure of the

economic model. In situations where a lower bound on welfare is useful (e.g., because even the lower bound estimate exceeds costs), they can provide sufficient information robust to some econometric problems. In other cases, they should be interpreted with caution. In simulations, the lower bound estimate is on the order of three-quarters of the true equivalent surplus.

2. Hedonic Capitalization Effects

2.1 The Hedonic Model

Let \mathcal{H} denote the set of houses in a region with typical element h and let \mathcal{J} 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. (Below, I relax this assumption and show it does not effect 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, but we will treat owner-occupied houses as if the household is rents from itself.

At any point in time t , households differ by their income y and by their current-period preferences, which can be represented by a twice differentiable quasi-concave conditional indirect utility function $v_i^t(y_i^t - p_h, g_h, \mathbf{x}_h)$, with $\partial v_i^t / \partial y_i^t > 0$. On the supply side of the market, the profit function for house h is $\pi_h = p_h - c_h(\mathbf{x}_h)$, where the cost function $c_h(\cdot)$ is twice differentiable. For simplicity, assume $c_h(\cdot)$ 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 equilibrium price function: $p_h^t = p^t(g_h^t, \mathbf{x}_h^t)$. The time superscript on the hedonic price function indicates that equilibrium hedonic prices may shift. In principle, these shifts may transpire from changes in the distribution of g , changes in household demands, or other changes in the economic environment. In each period, households maximize utility over a continuous choice set defined by the continuously differentiable hedonic 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

² This assumption continues to underlie majority of work on hedonic markets (e.g. Bajari and Benkard 2005, Bishop and Timmins 2019, Ekeland, Heckman, and Nesheim 2004, Heckman, Matzkin, and Nesheim 2010) as well as structural sorting models of locational choice (e.g. Bayer, Ferreira, and McMillan 2007, Kuminoff 2012, and Sieg et al.

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:

$$(2) \quad \frac{\partial v_i^t / \partial g}{\partial v_i^t / \partial y} = \frac{\partial p^t}{\partial g}.$$

Equation (2) is the standard tangency condition, with the derivative of the hedonic function with respect to an amenity is equal to marginal willingness to pay for the amenity at the optimal point.

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

$$(3) \quad \frac{\partial c_h}{\partial x_r} = \frac{\partial p^t}{\partial x_r}.$$

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

The basic problem is (i) to make inferences about non-marginal welfare effects from these primitive conditions and (ii) to estimate even the primitives when g or \mathbf{x} are endogenous, perhaps because of omitted variables.

2.2 Defining Capitalization Effects

To overcome the second issue, related to estimation, researchers have applied DD and related quasi-experimental approaches to the hedonic econometric model to identify the effects of exogenous changes in g (e.g. Bento et al. 2015, Cellini et al. 2010, Chay and Greenstone 2005; Currie et al. 2015; Davis 2004, 2011; Figlio and Lucas 2004; Greenstone and Gallagher 2008; Lang 2017, Linden and Rockoff 2008, Mastromonaco 2015, and Pope 2008). While effectively addressing an econometric problem, these methods have raised other questions about what precisely they identify. In an important advance in the literature, Klaiber and Smith (2013) and Kuminoff and Pope (2014) point out that if the hedonic price function shifts (either because of the changes in g itself or other changes in the economic environment), standard DD approaches conflate parameters describing two different equilibrium price functions.

2004). However recent work increasingly is considering 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 labor literature has long considered such dynamic optimization (e.g. Keane and Wolpin 1997).

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. From the perspective of the program evaluation literature, this can be viewed as a violation of SUTVA, particularly 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 treating some units has an "indirect effect" even on untreated units plus an additional "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 made by Hudgens and Halloran (2008), Tchetgen Tchetgen and VanderWeele (2012), VanderWeele and Tchetgen Tchetgen (2011), and others to consider effects defined by an entire policy—that is, by a change in g anywhere.

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 $\mathbf{x}_h^a(g_h^a)$ be the value of \mathbf{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 $t=1$.

We incorporate the violation of SUTVA by allowing for different potential prices at house h based not only on the value of g_h , but also based on the entire policy vector \mathbf{g} . The potential outcome for house h were we in the counterfactual state (with no houses treated), would be $p_h^1(\mathbf{g}_{-h}^{a'}, g_h^{a'}, \mathbf{x}_h^{a'}(g_h^{a'}))$. This price is illustrated in Figure 1, by price p_A , from the counterfactual price function evaluated at $g_h^{a'}$. The potential outcome for house h , if the rest of the market were under policy a^* , can be written as $p_h^1(\mathbf{g}_{-h}^{a^*}, g_h^{a'}, \mathbf{x}_h^{a'}(g_h^{a'}))$ if house h were not treated (receiving its counterfactual $g_h^{a'}$) and as $p_h^1(\mathbf{g}_{-h}^{a^*}, g_h^{a^*}, \mathbf{x}_h^{a^*}(g_h^{a^*}))$ if house h were treated (exogenously 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 counter-

factual, a' . Two distinctions here should be emphasized about this comparison in our setting relative to common assumptions. 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 between $t=0$ and $t=1$ 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, $p_B = p_h^1(\mathbf{g}_{-h}^{a^*}, \mathbf{g}_h^{a'}, \mathbf{x}_h^{a^*}(\mathbf{g}_h^{a'}))$ is not necessarily equal to $p_A = p_h^1(\mathbf{g}_{-h}^{a'}, \mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}(\mathbf{g}_h^{a'}))$ because the treatments at the other houses effect 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

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

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

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

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

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

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 $p_C - p_B$.

In theory, changing the treatment status of just house h by itself, as envisioned in the definition of DE, could itself have general equilibrium effects. Assumption A1 rules out such effects:

ASSUMPTION A1 (Local Non-interference). Let \mathbf{p}_{-h} be the vector of prices for all houses except h . Assume $\mathbf{p}_{-h}^1(\mathbf{g}_{-h}^{a^*}, \mathbf{g}_h^{a^*}, \mathbf{x}_h^1) \approx \mathbf{p}_{-h}^1(\mathbf{g}_{-h}^{a^*}, \mathbf{g}_h^{a'}, \mathbf{x}_h^1)$ for all h and all a^*, a' .

In words, changing the treatment status of only one house does not appreciably affect the price of any other house. This rules out, for example, "tipping" away from an unstable equilibrium. The local non-interference assumption can be taken to be a minimal instance of Palmquist's (1992) localized externality. Under this assumption, the direct effects DE can be interpreted as the movement along a constant hedonic price function, specifically the one prevailing in scenario a^* , from an untreated to a treated state, at a fixed \mathbf{x} .³ As the price function no longer depends on whether h alone was actually treated, with the local non-interference assumption, we can write these potential effects more simply as

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

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

That is, the evaluation of $p_h^a(\mathbf{g}_h^1, \mathbf{x}_h^1)$ need not account for the general equilibrium price effects of changing the treatment status of only house h .

As written, TE, IE, and DE all include any effects mediated through changes in \mathbf{x} .⁴ For example, improvements in g might motivate households to improve the house in other (observable) ways, or it might trigger resorting with new households wanting to change the house. Variants of these treatment effects that net out the portions mediated through changes in \mathbf{x} can be defined for all three. Define the *total unmediated effect* (TUE) and the *direct unmediated effect* (DUE) at $\tilde{\mathbf{x}}_h$ respectively by:

$$\text{TUE}_h(a^*) = p_h^{a^*}(\mathbf{g}_h^{a^*}, \mathbf{x}_h^{a^*} = \tilde{\mathbf{x}}_h) - p_h^{a'}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'} = \tilde{\mathbf{x}}_h)$$

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

³ To avoid the general equilibrium effects of changing the treatment for one unit, Hudgens and Halloran (2008) impose a particular randomization assumption that under any policy a , the number of treated units is fixed. Thus, \mathbf{g}_{-h}^{1,a^*} is conditioned on the value of \mathbf{g}_h^1 in their definition of the direct effect. To avoid this awkward construction, VanderWeele and Tchetgen (2011) propose an alternative definition of the direct effect in which \mathbf{g}_{-h}^{1,a^*} is fixed, but which no longer decomposes the total effect. The local non-interference assumption provides an alternative way to address this issue. Under this assumption, both definitions of the direct effect are equivalent.

⁴ As noted by Tchetgen Tchetgen and VanderWeele (2012), terms like "direct" and "indirect" can be somewhat confusing in the presence of both interference and mediation. I use "indirect" to mean changes in the hedonic price function (as defined previously) and "mediated" to be the effect through changes in \mathbf{x} attributable to changes in g .

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

Each of the individual effects defined above have their respective group averages. Define the average total and direct unmediated effects, respectively, as follows:

$$\overline{\text{TUE}(a^*)} = \frac{1}{H} \sum_h \text{TUE}_h(a^*),$$

$$\overline{\text{DUE}(a^*)} = \frac{1}{H} \sum_h \text{DUE}_h(a^*).$$

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

$$\overline{\text{DUE}_{\mathcal{H}'}(a^*)} = \frac{1}{\sum_h 1(h \in \mathcal{H}')} \sum_{h \in \mathcal{H}'} \text{DUE}_h(a^*).$$

When we have a particular treatment program in mind, a particular special case of $\overline{\text{DUE}_{\mathcal{H}'}(a^*)}$ is where \mathcal{H}' is simply the subset of treated houses. This is the average direct unmediated effect on the treated, or $\overline{\text{DUET}(a^*)}$.

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 effects."⁵ $\overline{\text{TE}(a^*)}$ 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

⁵ Linking hedonics to "capitalization" has a long history. When Frederick Waugh (1929) first used hedonic methods to explain vegetable prices as a function of their attributes, Eveline Burns commented that the effect was analogous to Ricardian rent for cross-section differences in the fertility of land. Early interpretations of capitalization in time include Lind (1973) and Starrett (1981). By the 1980s, "capitalization" was used interchangeably in potentially confusing ways to refer to cross sectional capitalization or to capitalization over time.

and Smith (2013) have rightly emphasized, $\overline{\text{TUE}(a^*)}$ is not the average willingness to pay for program a^* . The problem is that it conflates the direct and the indirect unmediated effects. Indeed, it is hard to give it any welfare interpretation except in the special case where the hedonic function does not in fact change, in which case the results of Palmquist (1992) apply. In that case, there is no indirect effect, so the total and direct effects coincide.

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.⁶ 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). Here, the figure shows that g does not change from the baseline, but that assumption is not necessary.⁷ The treated units have an additional shock to g and end up at p_C . Thus, the identified effect from a DD comparison is $(P_C - P_A) - (P_B - P_A) = (P_C - P_B)$, which is the movement along the ex post hedonic price function from treatment, or DUE. As shown in Section 4, this concept can be interpreted as a lower bound on welfare.

3.1 Identification and Estimation of Capitalization Effects: The Linear Case

Consistently with most hedonic work, let us first develop the argument with a linear model. In-

⁶ 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 $p_h^1(\mathbf{g}_{-h}^{1,a'}, g_h^{1,a'})$ in the expression for TE. In effect, this assumption is that $p_h^{a'}$ is observed, as p_h^0 . Another is that observations are available at other markets that are not treated and that between-market trends are assumed to be such that identification can leverage inter-city comparisons. See Hudgens and Halloran (2008) and Manski (2013). Crépon et al. (2013) illustrate the idea.

⁷ In principle, one could allow g to change for the untreated, so long as it would have changed in the same way for the treated but for the treatment. For example, Chay and Greenstone (2005) and Bento et al. (2015) consider the price effects of air quality improvements triggered by regulatory thresholds crossed in some counties but not others. One can identify the direct effects of the change in air quality triggered by those regulations, while still allowing for changing air quality nationally.

roducing error terms into the price functions, for any individual house h , the ex ante and a^* hedonic price functions and their difference are, respectively:

$$(4a) \quad p_h^0 = \alpha^0 + \beta^0 g_h^0 + \gamma^0 \mathbf{x}_h^0 + \xi_h + \varepsilon_h^0,$$

$$(4b) \quad p_h^{a^*} = \alpha^{a^*} + \beta^{a^*} g_h^{a^*} + \gamma^{a^*} \mathbf{x}_h^{a^*} + \xi_h + \varepsilon_h^{a^*},$$

$$(4c) \quad dp_h^{a^*} = d\alpha^{a^*} + d\beta^{a^*} g_h^0 + \beta^{a^*} dg_h^{a^*} + d\gamma^{a^*} \mathbf{x}_h^0 + \gamma^{a^*} d\mathbf{x}_h^{a^*} + d\varepsilon_h^{a^*},$$

where in Equation (4c) the differences are taken from the baseline, not from the unobserved counterfactual. This equation 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 (4b), as the parameters are independent of the value of $g_h^{a^*}$. If the hedonic price function does not change between (4a) and (4b), as with Palmquist's (1992) localized externality, then we can suppress time-scenario superscripts in the parameters and Equation (4c) collapses to

$$(5) \quad dp_h^{a^*} = \beta dg_h^{a^*} + \gamma d\mathbf{x}_h^{a^*} + d\varepsilon_h^{a^*}.$$

In this case, it is clear that 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, in the general case where the hedonic function does shift, the true model is (4c), so (5) of course is mis-specified. Consequently, it suffers from omitted variable bias: g^0 and \mathbf{x}^0 belong in the model but are omitted. In their recent discussion of capitalization, Kuminoff and Pope (2014) refer to this problem as "conflation bias," as the estimates conflate marginal willingness to pay at a point in time (i.e. β^0 or β^{a^*}) with changes in the hedonic price function. As they show, conflation bias is an example of omitted variable bias. Clearly, if g^0 and \mathbf{x}^0 are included in the model, as in Equation (4c), the linear model potentially can identify β^{a^*} , the ex-post marginal willingness to pay under the scenario. Thus, any flaw in the model arises from failure to properly condition on baseline observables, not with the economic logic of differencing prices from two equilibria *per se*.

Of course, including g^0 and \mathbf{x}^0 in a linear model, as in Equation (4c), may well raise additional estimation issues. In particular, although it allows for the existence of an unobserved time invariant effect, ξ_h , Equations (4) still require a conditional zero mean assumption on $d\varepsilon$ to estimate the full set of parameters in (4c) from OLS. Unfortunately, $d\varepsilon$ may well be correlated with g^0 : for example, houses near high levels of pollution may be depreciating in unobserved ways. However,

important (if incomplete) information can still be identified under a weaker conditional independence assumption, in which dg^{a^*} is independent of $d\varepsilon^{a^*}$ conditional on g^0 and the other observables: $(d\varepsilon \perp dg \mid g^0, \mathbf{x}^0, d\mathbf{x})$. Such conditional independence would allow identification of β^{a^*} , even if β^0 were biased.

The fact that it is the ex post hedonic price parameter under the realized scenario, 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: $g^0=0, g^{a^*} \in \{0,1\}$. Examples might include cleanup of Superfund sites (Gamper-Rabindran and Timmins 2013, Greenstone and Gallagher 2008), discovery of a cancer cluster (Davis 2004), closing of large polluting facilities (Currie et al. 2015, Mastro Monaco 2015), arrival of a sex offender (Linden and Rockoff 2008) and so forth.⁸ Relaxing the linearity inherent in Equations (4), we can define the potential outcomes using the following semi-parametric assumptions:

$$(6a) \quad p_h^0 = \boldsymbol{\gamma}^0 \mathbf{x}_h^0 + \varepsilon_h^0,$$

$$(6b) \quad p^{a^*}(g_h^{a^*} = 0) = \boldsymbol{\gamma}_{g^1=0}^{a^*} \mathbf{x}_h^{a^*} + \varepsilon_{g^{a^*}=0,h}^{a^*},$$

$$(6c) \quad p^{a^*}(g_h^{a^*} = 1) = \boldsymbol{\gamma}_{g^1=1}^{a^*} \mathbf{x}_h^{a^*} + \varepsilon_{g^{a^*}=1,h}^{a^*},$$

where the $\boldsymbol{\gamma}$ vectors include an intercept term. Equation (6b) represents houses that are not treated ex post, Equation (6c) 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).⁹ It again implicitly relies on the local non-interference assumption A1,

⁸ The model of this sub-section could be extended to include multi-valued or even continuous treatments along the lines suggested by Imbens (2000) and Hirano and Imbens (2004).

⁹ Note that $E[\varepsilon_{g^t,h}^{a^*} \mid \mathbf{x}]$ need not be zero. If we could observe $[\varepsilon_{g^1=1,h}^{a^*} \mid \mathbf{x}]$ and $E[\varepsilon_{g^1=0,h}^{a^*} \mid \mathbf{x}]$ for the same house, we could take the difference as our estimate of the effect of g conditional on \mathbf{x} . Of course, we don't observe the latter for treated houses, which is the standard missing counterfactuals problem.

as $\boldsymbol{\gamma}_g^{a^*}$ 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):

$$(7) \quad E[\varepsilon_{g^1=0}^{a^*} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 1] = E[\varepsilon_{g^1=0}^{a^*} - \varepsilon^0 \mid \mathbf{x}^0, g^{a^*} = 0].$$

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, as well as the usual requirement of overlapping support, a conditional DD estimand can identify the average direct unmediated effect on the treated. This is stated more formally in the following lemma.

LEMMA. Given A1, A2, and the model of Equations (6),

$$(8) \quad \begin{aligned} & E \left[\left(p^{a^*}(g^{a^*} = 1) - \boldsymbol{\gamma}_{g^1=0}^{a^*} \mathbf{x}^{a^*} \right) - \left(p^0 - \boldsymbol{\gamma}_{g^1=0}^{a^*} \mathbf{x}^0 \right) \mid \mathbf{x}^0, g^{a^*} = 1 \right] - \\ & E \left[\left(p^{a^*}(g^1 = 0) - \boldsymbol{\gamma}_{g^1=0}^{a^*} \mathbf{x}^{a^*} \right) - \left(p^0 - \boldsymbol{\gamma}_{g^1=0}^{a^*} \mathbf{x}^0 \right) \mid \mathbf{x}^0, g^{a^*} = 0 \right] \\ & = E \left[\left(\left(p^{a^*}(g^{a^*} = 1) - p^{a^*}(g^1 = 0) \right) \mid \mathbf{x}^0, g^{a^*} = 1 \right) - \boldsymbol{\gamma}_{g^1=0}^{a^*} \left((d\mathbf{x} \mid \mathbf{x}^0, g^{a^*} = 1) - (d\mathbf{x} \mid \mathbf{x}^0, g^{a^*} = 0) \right) \right] \\ & = \overline{\text{DUET}(a^*)}. \end{aligned}$$

Proof: The first equality follows immediately from Equations (6) and Assumption A2. The second follows from Equations (6) and Assumption A1. The third follows from the definition of DUET.

$\overline{\text{DUET}(a^*)}$ might be estimated using a regression-adjusted difference-in-differences 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 \mathbf{x} . The estimand defined in the Lemma recovers an analogous effect, for those houses

actually treated by the policy, but using weaker econometric assumptions. It is the effect of moving along the ex post hedonic, for treated houses.

3.3. Quasi-Experimental Designs in Practice

Section 4 discusses the economic interpretation of this estimand. But before turning to that discussion, it may be useful to consider other quasi-experimental designs used in the literature.

Repeated cross sections

The analysis and discussion so far have been in terms of panel data at the individual level, with repeated observations for each h . In the context of housing, this is known as the "repeat sales" model. However, the basic insights hold if we only observe repeated cross sections, with a panel at the group level. Suppose, for example, that individual houses can be grouped into local communities c , with ξ defined at the community level, but \mathbf{x} and maybe g still observed at the individual level. Then, to take the linear case, we can rewrite model (4a)-(4c) as

$$(9a) \quad p_h^0 = \alpha^0 + \beta^0 g_h^0 + \gamma^0 \mathbf{x}_h^0 + \xi_c + \varepsilon_h^0,$$

$$(9b) \quad p_h^{a^*} = \alpha^{a^*} + \beta^{a^*} g_h^{a^*} + \gamma^{a^*} \mathbf{x}_h^{a^*} + \xi_c + \varepsilon_h^{a^*}.$$

$$(9c) \quad p_h^t = \alpha^0(1-D^t) + \alpha^{a^*}D^t + \beta^0 g_h^0(1-D^t) + \beta^{a^*} g_h^{a^*}D^t + \gamma^0 \mathbf{x}_h^0(1-D^t) \\ + \gamma^{a^*} \mathbf{x}_h^{a^*}D^t + \xi_c + \varepsilon_h^t,$$

where D^t is a dummy variable for the second period. Here, Equation (9c) simply stacks Equations (9a) and (9b). Thus, the model still can be estimated with individual level data on \mathbf{x} and community fixed effects. While some studies have used repeat sales (e.g. Figlio and Lucas 2004), most hedonic DD designs are based on this kind of panel model. Examples include Davis (2004), Linden and Rockoff (2008), Mastro Monaco (2015), Muehlenbachs et al. 2015, and Pope (2008). In principle, one also could allow the spatial dummies to be time-varying. Additionally, for the linear case, we can always collapse the data to the community-year level, with \mathbf{x} now observed as the group-year average. Consequently, analysts can estimate these DD models with repeated cross sections with group-level fixed effects, or even group level means (i.e. they can use either the within or the between estimator) (e.g. Currie et al. 2015, Lucas 2011).

Omitting baseline characteristics in Equation (4c) (as in Equation 5) is equivalent to im-

posing the condition that $\beta^{a^*} = \beta^0$ in Equation (9c).¹⁰ Most previous studies do impose this assumption of a constant slope for g over time. As discussed above, this restriction likely introduces conflation bias, yielding neither β^{a^*} nor β^0 , unless the changes are small and over a short time period (Kuminoff and Pope 2014). Interesting exceptions to this rule are Davis (2004) and Figlio & Lucas (2004), both of which look at information treatments (new evidence of a cancer cluster and new signals about school quality, respectively). In both studies, there is no information prior to the treatment so there are no treated observations in the ex ante period, while fixed effects control for time-invariant differences at the geographic scale of the treatment. With no variation in g_h^0 , imposing $\beta^{a^*} = \beta^0$ cannot bias the estimates of β^{1,a^*} (though ideally one would still allow γ to vary over time).

Ultimately, recovering β^{a^*} requires an exogenous source of variation in dg . The better our understanding of the variation in dg , and why it might be orthogonal to $d\varepsilon$ conditional on (g^0, \mathbf{x}^t) , the closer fixed effects designs approach quasi-experimental methods. In the case of Davis (2004), Figlio and Lucas (2004), and Mastromonaco (2015), the shock to g and the hedonic equilibrium caused by new information is plausibly exogenous, either because of chance discovery or new disclosure rules imposed at the state or national level. These designs allow for the recover of β^{a^*} . Similarly, the arrival of a sex offender to a very local neighborhood, studied by Linden and Rockoff (2008) and Pope (2008), is plausibly random after conditioning on community fixed effects. Allowing β and γ to vary over time in these studies would have identified $\overline{\text{DUET}(a^*)}$.

In their study of the value of brownfield remediation, Haninger et al. (2017) use an alternative DD strategy in which they use only cross-sectional variation, comparing the difference between housing prices within 0.2 km of cleaned up brownfields and those within 0.2 km of remaining brownfields to the analogous difference for houses within 0.5 km of brownfields. They argue this strategy enables them to identify the ex post hedonic function. Currie et al. (2015) use a similar design for openings and closings of Toxic Release Inventory (TRI) facilities, but extend the logic to *triple* differences. Essentially, they find the difference-in-differences between property values

¹⁰ To see the connection, note that we can rewrite the terms $\beta^0 g_c^0 (1-D^l) + \beta^{a^*} g_c^{a^*} D^l + \xi_c$ by adding and subtracting $\beta^0 g_c^0 D^l$ and $\beta^{a^*} g_c^0 D^l$. Grouping terms, this becomes $(\beta^0 g_c^0 + \xi_c) + (\beta^{a^*} g_c^{a^*} - \beta^{a^*} g_c^0 + \beta^{a^*} g_c^0 - \beta^0 g_c^0) D^l = (\beta^0 g_c^0 + \xi_c) + (\beta^{a^*} dg_c^0 + d\beta g_c^0) D^l$. That is, it is equivalent to entering g_c^0 for both time periods (which then is absorbed by the fixed effect) and identifying β^{a^*} from dg_c and $d\beta$ from g_c^0 using just the second period observations, as we would if first differencing (only now there are more than one observation per community). Imposing constant slopes, i.e. $d\beta = 0$, is equivalent to dropping g_c^0 , as discussed above.

within 1 mile of a site that opened (or closed) relative to one that did not change status, and compare that difference-in-difference for its matched control between 1-2 miles from the site. Though they have no quasi-experimental variation in openings and closings per se, this strategy relies on the fairly weak assumption that openings and closings may be correlated with unobserved economic conditions within 2 miles, but not differentially between 1 and 2 miles.

Regression discontinuity designs

The exogeneity of g can be made more credible using regression discontinuity (RD) designs (e.g. Baum-Snow and Marion 2009, Cellini et al. 2010, Chay and Greenstone 2005, Gamper-Rabindran and Timmins 2013, Greenstone and Gallagher 2008, and Lang 2018). With RD, the key source of variation in g is the movement across a threshold of some continuous forcing variable, with unobservables varying continuously across the threshold. Within the neighborhood of the threshold, assignment to the treatment can be thought of as being as good as randomly assigned, as in a controlled experiment (Lee and Lemieux 2010). This random assignment has two key implications for interpreting hedonics. First, the average treatment effect is the same as the treatment on the treated, so RD designs can identify the latter. Second, RD estimates are unbiased even when researchers do not include a full set of controls, though they are less efficient in small samples.

In principle, RD designs can be used in either a panel or a cross-sectional setting. Consider the panel-data setting first. As an example, Chay and Greenstone (2005), estimating hedonic regressions of housing prices on air quality, persuasively argue that recessions or local economic shocks can simultaneously reduce housing prices in unobserved ways while improving air quality (by dampening economic activity), thus biasing DD (or fixed effects) hedonic estimates of air quality downward. They argue that national ambient air quality thresholds, which trigger regulation, are a plausible source of exogenous variation in air quality. Accordingly, they regress 1970-1980 changes in housing prices on 1970-1980 changes in air quality, using mid-decade regulatory thresholds as an instrument. They essentially estimate Equation (5), using the RD as a source of variation in dg . Although they do not control for baseline conditions, the as-good-as-random argument for RD suggests this is not necessary (though it may be efficient). As discussed by Kuminoff and Pope (2014), if areas just below the threshold are a valid counterfactual to areas just above (with similar baseline conditions and similar trends in prices), this RD strategy implicitly conditions on g^0 . Thus, it identifies $\overline{\text{DUET}(a^*)}$.

Similarly, Greenstone and Gallagher (2008) argue that a discontinuity in priority scores assigned to Superfund's National Priority List is an exogenous source of variation in the assignment of cleanup to toxic sites (see also Gamper-Rabindran and Timmins 2013). However, in this case, they model the ex post cross-sectional hedonic price function directly, regressing ex post prices on cleanup status (instrumented by the discontinuity in scores), controlling for baseline conditions. Thus, again, this strategy identifies the movement along the ex post hedonic induced by the treatment.

Finally, Cellini et al. (2010) and Lang (2018) consider the dynamic effects over time of funding school capital projects and of conservation of open space, respectively. They use discontinuities in the voting outcome (e.g. at 50%) as an exogenous source of variation in passage of referenda authorizing such funding. Both papers use longer time periods and innovatively allow the referenda to have evolving effects over time. However, they also include referenda that pass at different points in time. Thus, they implicitly impose the condition that $\beta^{a^*} = \beta^0$ for each lag. As they are simultaneously identifying the effects of multiple cross sections, the usual RD arguments do not apply to the errors caused by mis-specifying an evolving hedonic function.

In sum, RD designs can identify the direct treatment effect on the treated. However, they only identify such effects locally, around the threshold of treatment. Yet it is the average direct effect on *all* the treated that is usually of interest, and which will be shown in Section 4 to be a lower bound on welfare. Two responses to this limitation are possible. One is simply to assume uniform treatment effects for all values of the forcing variable (Greenstone and Gallagher 2008). As an alternative, Dong and Lewbel (2015) have proposed using *kinks* in the hedonic price function around the threshold to identify changes in slopes of the running variable, and to extrapolate these changes out over the domain of the treated. They discuss and evaluate this approach in the context of Greenstone and Gallagher's hedonic study.

Instrumental variables designs

Even in settings where an RD design is not available, one may not want to impose the conditional independence assumptions required for fixed effects regression or DD matching. If so, one could invoke additional exclusion restrictions and use instrumental variables (IV). For example, adapting the work of Chay and Greenstone (2005), Bento et al. (2015) recently have used the *proportion* of years in a time window that exceed the regulatory threshold as an instrument for air quality,

making it continuous rather than discrete. In a very different application, Gopalakrishnan et al. (2011) use coastal geological features to instrument for beach width in a cross-sectional hedonic price regression of coastal properties. As with RD, the use of IV again raises the question of local average effects, and if interpreted through the lens of this paper would seem to require an assumption of homogenous effects. It also raises the possibility of identifying effects in multiple cross-sections, a point I return to at the end of the paper.

4. The Welfare Interpretation of Capitalization Effects: The Direct Unmediated Effect is a Lower Bound on Hicksian Equivalent Surplus

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 value of a house at the new and old level of the amenity respectively, netting out effects mediated through dx : $p^{a^*}(g^1, \mathbf{x}^0) - p^{a^*}(g^0, \mathbf{x}^0) = \int_{g^0}^{g^1} \frac{\partial p^{a^*}(g, \mathbf{x}^0)}{\partial g} dg$. The most important thing to emphasize is that the effect is based on the ex post hedonic function. In other words, the counterfactual is *not* the price in the ex ante scenario; nor is it what the price of the house would have been in the absence of the policy, because that counterfactual equilibrium is never observed. The counterfactual for a treated observation is what the price of that one individual house would have been if its g were not affected by the policy but the policy had otherwise gone forward (and the price function had thus shifted). The different counterfactuals are only the same in the case of Palmquist's (1992) localized externality; in general they are different.

What is the economic interpretation of this estimand? Summed over houses, the expression $\int_{g^0}^{g^1} \frac{\partial p^{a^*}}{\partial g} dg$ is a lower bound on the sum of the consumer's Hicksian equivalent surplus (ES) for the improvement plus the change in profits, for *all* households in the city, not just those living (or owning) the treated houses at some point of time. The 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. 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).

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 \mathbf{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. This can be seen immediately in Figure 2. The figure depicts a Marshallian bid function for g intersecting the two marginal price functions at g^0 and g^1 . It also depicts a Hicksian demand function evaluated at u^1 , which intersects the Marshallian function at g^1 . Whereas Hicksian equivalent variation would be the area under this Hicksian demand function from the level of g chosen under p^0 to achieve u^1 , depicted here as $g(p^0(\cdot), u^1)$, up to g^1 , Hicksian equivalent surplus is the area under the function from g^0 to g^1 . The difference is in the point of evaluation. The former is associated with the solutions to the expenditure minimization problem given the two hedonic price functions and u^1 . The latter is the Hicksian value for the realized change in g , determined from the Marshallian problem. Clearly, $\int_{g^0}^{g^{a^*}} \frac{\partial p^{a^*}}{dg} dg$ is a lower bound on $ES = \int_{g^0}^{g^{a^*}} h(g, u^1) dg$. The argument is analogous to the well-known fact that a Paasche quantity index is a lower bound for the value of a quantity change.

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 \mathbf{x} (for example, with upgrades or additions occurring to the housing stock in response to the policy), and resorting of households. The required assumption on profits is:

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

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

A special case of Assumption A3 would be the case of zero profits. However, the weaker assumption of non-negative profits is all that is required. This assumption seems reasonable intuitively: landlords would not make the investment in changes to \mathbf{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, $H^T \cdot \overline{\text{DUET}(a^*)} \leq (\text{ES} + \Delta \text{profits})$ 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 the levels of \mathbf{x} over time, but does not include these in the welfare measure.

Proof: See Appendix.

The proposition states that $\overline{\text{DUET}(a^*)}$, times the number of treated units, is a lower bound on welfare for *all* households effected by the policy, whether directly or indirectly. All households are included because they potentially are affected by the policy through the changing price function.¹¹ Note that, although panel data are used to control econometrically for unobservables, $\overline{\text{DUET}(a^*)}$ 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 economic concept. If there are changes in current period demands or tastes, then the result remains valid, but the evaluation is from the perspective of ex post preferences.

The formal proof follows the outline of the verbal argument in Bartik (1988), clarifying a few ambiguous points.¹² Denote the expenditure function for household i as $e_i(p(\cdot), u)$ where $p(\cdot)$ is the hedonic price function and the price of other goods is normalized to one. It is the solution to the expenditure minimization problem when the household faces hedonic price function $p(\cdot)$. Denote the restricted expenditure function as $\tilde{e}_i(p(g, \mathbf{x}), g, \mathbf{x}, u)$; it is the solution to the expenditure minimization problem when the household is constrained to choose the bundle (g, \mathbf{x}) . The basic idea is to define the welfare measure as follows:

¹¹ As noted above, the bound is exact when the price function does not change. In that case, the welfare effect also applies only to landlords of the treated houses (see Palmquist 1992).

¹² 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 proof.

Additionally Kanemoto (1988) provided another related bounding proof which has a somewhat similar flavor to Bartik's. 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.

$$(10) \quad dW = \sum_i [\tilde{e}_i(p^{a'}(g_{i(a')}, \mathbf{x}_{i(a')}), g_{i(a')}, \mathbf{x}_{i(a')}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_h [(p^{a*}(g_h^{a*}, \mathbf{x}_h^{a*}) - p^{a'}(g_h^{a'}, \mathbf{x}_h^{a'})) - (c(g_h^{a'}, \mathbf{x}_h^{a*}) - c(g_h^{a'}, \mathbf{x}_h^{a'}))].$$

where $(g_{i(a')}, \mathbf{x}_{i(a')})$ 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, the Hicksian equivalent surplus (ES) for the change in g . If preferences have changed over time, then this measure reflects preferences for $t=1$. The second term in brackets is the change in landlord profits. It is the change in rents, resulting from both the shift in the hedonic price function and adjustments in \mathbf{x} as well as exogenous changes in g , minus the change in costs, evaluated at baseline levels of g .

As shown in the appendix, using Assumption A.3, this measure is equivalent to:

$$(11) \quad dW = \sum_i [\tilde{e}_i(p^{a*}(g_{i(a')}, \mathbf{x}_{i(a')}), g_{i(a')}, \mathbf{x}_{i(a')}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_h [(p^{a*}(g_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(g_h^{a'}, \mathbf{x}_h^{a'})) - (c(g_h^{a'}, \mathbf{x}_h^{a*}) - c(g_h^{a'}, \mathbf{x}_h^{a'}))] \\ + \sum_h (p^{a*}(g_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(g_h^{a'}, \mathbf{x}_h^{a*})).$$

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. Since 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 price 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. In simulations presented below, the lower bound welfare estimate for improvements in g ranges from 92% of ES for small changes to as low as low as 75% of ES for larger changes.

5. Extensions and Discussion

5.1. 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.

5.2. 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 4.1 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 (10), 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 $(g_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'})$, 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 $(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')})$.

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 (10), we replace the expression

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

with the expression

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

The term $e_i^{TC(R)}(p^{a^*}(), u_i^{TC(R),a^*})$ 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 $\tilde{e}_i^{TC(a^*)}(p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}, u_i^{TC(R),a^*})$ 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 appendix, all other terms also remain unchanged, so the argument flows through without any other alterations. We are left in the end with the expression:

$$(12) \quad \begin{aligned} dW = & \sum_i [\tilde{e}_i^{TC(a^*)}(p^{a^*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}, u_i^{TC(R),a^*}) - e_i^{TC(R)}(p^{a^*}(), u_i^{TC(R),a^*})] \\ & + \sum_h [(p^{a^*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a^*}) - p^{a^*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) - (c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a^*}) - c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}))] \\ & + \sum_h (p^{a^*}(\mathbf{g}_h^{a^*}, \mathbf{x}_h^{a^*}) - p^{a^*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a^*})). \end{aligned}$$

If it were not for the transactions costs paid, this expression would be the same as Equation (11) evaluated at a different utility level. Consequently, most of the argument flows through. However, we have to account for the fact that the reference points for the two expenditure functions in (12) are different: one involves $TC(a^*)$ whereas the other involves $TC(R)$. In the appendix, I show that 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, $H^T \cdot \overline{DUE_{\mathcal{H}'}(a^*)} \leq (ES^{TC} + \Delta\text{profits})$ 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 x over time, but does not include these in the welfare measure.

Proof: See the Appendix

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

5.3 The Direct Mediated Effect.

The analysis to this point has 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? As noted above, it should not matter in an RD design, at least in expectation. In a DD design, unfortunately, the observed measure incorporates the gross benefits of general equilibrium adjustments to x , without netting out the costs of providing them.¹³ This can undermine the lower bound interpretation, although we can still establish sufficient conditions under which it holds.

To see this, rewrite Equation (11) as:

¹³ An exception might be other neighborhood amenities that are freely provided but change as a result of the policy. For example, g might be transportation infrastructure, which could affect air quality. One would not want to net these effects out. It might be simplest in these cases to think of g as a vector.

$$(13) \quad dW = \sum_i [\tilde{e}_i(p^{a^*}(g_{i(a')}, \mathbf{x}_{i(a')}), g_{i(a')}, \mathbf{x}_{i(a')}, u_i^{a^*}) - e_i(p^{a^*}(\cdot), u_i^{a^*})] + \\ - \sum_h (c(g_h^{a'}, \mathbf{x}_h^{a'}) - c(g_h^{a'}, \mathbf{x}_h^{a'})) + \sum_h (p^{a^*}(g_h^{a'}, \mathbf{x}_h^{a'}) - p^{a^*}(g_h^{a'}, \mathbf{x}_h^{a'})).$$

Here, the first set of terms remains the same as in Equation (11), as do the change in costs, while the last set of terms is the direct effect gross of any change in \mathbf{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

$$(14) \quad \sum_h (c(g_h^{a'}, \mathbf{x}_h^{a'}) - c(g_h^{a'}, \mathbf{x}_h^{a'})) \leq \\ \sum_i [\tilde{e}_i(p^{a^*}(g_{i(a')}, \mathbf{x}_{i(a')}), g_{i(a')}, \mathbf{x}_{i(a')}, u_i^{a^*}) - e_i(p^{a^*}(\cdot), u_i^{a^*})].$$

Essentially, the change in costs from the change in \mathbf{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.

This analysis suggests it is important to control for changes in \mathbf{x} whenever possible to have the cleanest welfare interpretation. However, that interpretation will still be valid if changes in \mathbf{x} are small or if changes in the cost of providing them are negative. It is worth emphasizing here that the question centers on changes in \mathbf{x} , not levels. Unobserved, unchanging levels of \mathbf{x} cancel out in the comparison, which motivate the use of DD strategies at the outset.

5.4 Changes in Population Density or Housing Density

The discussion so far has focused on the prices of a fixed set of housing units. What if population mobility leads to systematic changes in population density and/or construction of new housing units? Consider first changes in population density over a fixed set of housing units. As presented, the model allows households to systematically re-adjust among housing units, so that population density in space may shift depending on households' composition. Furthermore, households may come and go from an open region, affecting density in the region. As discussed in Section 5.1, such affects may change true welfare, but they do not change the construction of the lower bound on welfare, which depends only on the price effects of treated housing units.

Changes in the number of housing units on existing parcels can be accommodated by treating the parcel as the level of observation, h , rather than the housing unit. Then changes in housing units on the parcel could be thought of as changes in \mathbf{x}_h and incorporated into the analysis like

other changes in \mathbf{x} . As discussed in Section 5.3, it is generally best to condition on such changes, to net out their price effects. Finally, there remains the possibility of subdividing or assembling parcels. Future research might consider how to incorporate such reconfigurations into the model.

5.5 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}' . That is, given the equilibrium a^* , we might want 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 $\overline{\text{DUE}}_{\mathcal{H}'}(a^*)$.¹⁴

Under the appropriate assumptions, then, one could claim to identify the movement along the observed ex post hedonic function $p^{a^*}(\cdot)$ 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 $p^*(\cdot)$.

This movement along $p^{a^*}(\cdot)$, 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 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.

¹⁴ Specifically, the assumptions required would be:

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

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

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).

But another way the set of treated houses \mathcal{H}' may differ is to imagine an alternative counterfactual 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.

6. 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:

$$(15) \quad v_i(p, g_j, y_i) = \frac{y_i + \gamma g_j}{\exp\left(\frac{\exp(\alpha_i + \delta_i g_j) p^{1+\beta}}{1+\beta}\right)}$$

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:

$$(16) \quad x_i^* | g_j = \exp(\alpha_i + \delta_i g_j + \beta \ln p + \ln y).$$

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 (15), 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 $\gamma=100$. 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 1 shows the results from the main DD estimates (Equation 4c), 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 that in Equation (4c), 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 linear estimates of 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 hedonic can capture it well, whereas for larger changes the approximations fails to capture the curvature in the willingness to pay 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 can be 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 1 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 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, since 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 x . These models represent the direct mediated effect discussed in Section 5.4. 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 (14) is not satisfied here.

Table 2 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 1, with no transactions costs, for comparison. For the remaining rows, the true ES changes to reflect Equation (12) rather than (11), as discussed in Section 5.2. 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 1. 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 1 as well, again with an upward bias on the lower bound measure (these estimates are not reported but are available upon request).

7. 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 recently considered by Currie et al. (2015) and Mastromonaco (2015). In particular, I estimate a variant of Equation (9c), using two cross sections of individual houses and treating local geographic areas as the panel unit. 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 (9c). Thus, this approach identifies $\overline{DUE(a^*)}$ and the lower bound on welfare.

The specific application is to the Los Angeles area (including all of LA and Orange Counties and portions of Riverside, San Bernardino, and Ventura Counties), between 1995 and 2000. Data on toxic emissions come from the US EPA's TRI database. I consider only air emissions, either directly from stacks (point sources) or fugitive (nonpoint). 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 subject to measurement error, if they 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 between 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.)

The data include the latitude and longitude of facilities. Exposure to TRI facilities was imputed to a house in a given census block in two ways. The first uses the number of facilities within two miles. The second is a distance-weighted version, where a facility has a weight of $\max\{0, 1 - \frac{1}{2}Distance\}$, 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 given census block receives a weight of

zero, a facility 1 mile away is given a weight of $\frac{1}{2}$, and a facility co-located with a census block is given full weight. The top panel of Table 3 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, using either measure.

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.¹⁵ The bottom panel of Table 3 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 change over time. Additionally, they are based on population density, so they have widely varying sizes within a county. 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. Alternatively, I consider grids of $\frac{1}{2}$ -mile or 2-miles.

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

¹⁵ 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.

in exposure interacted with a time trend, plus contemporaneous exposure and the hedonic variables listed in Table 3. Table 4 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. Note, first, that the price levels are not monotonic by category. More importantly, the trends are fairly parallel—remarkably so for the last three years of the period. This evidence reassuringly suggests that TRI sites were not closing in areas that are already gentrifying.

For the main model of interest, I regress log price on nearby TRI sites, a cubic of lot area, living area, bathrooms, bedrooms, and age, with separate functions estimated for 1995 and 2000, plus time invariant fixed effects for the grid cell, and year-quarter dummies to pick up proportionate housing price inflation. This is the generalized DD regression of Equation 9c. After estimating the model, I use the 2000 price function to estimate the direct unmediated effect of the observed 1995-2000 change in exposure. I then add up the total value for the study area, re-weighting the data so as to reflect the number of owner-occupied housing units in each Census Tract.¹⁶ This is the lower bound on ES. I then bootstrap these values, which are a weighted forecast, to obtain confidence intervals on the lower bound welfare measure.

Table 5 presents the results using 1-mile grid cells for fixed effects. Each row represents a separate regression. The first column presents the coefficient on TRI exposure in 1995, which as discussed in Section 3.1 is only identified under the strong assumptions that changes in unobservables are uncorrelated with baseline conditions, and is not used in the construction of the lower

¹⁶ Specifically, I first predict the direct unmediated effect at each house in the 2000 transactions sample, then compute Census tract level means of these values, and finally aggregate up to a total market-wide value using the total number of owner-occupied homes in the tract.

bound on ES. The second column presents the coefficient for 2000, which is identified under the weaker assumption that changes in unobservables are uncorrelated with changes in exposure, conditional on baseline exposure and other variables. All standard errors and confidence intervals cluster by grid cell. The first row shows the results from using the number of TRI facilities within two miles as the measure of exposure, with 1-mile-grid fixed effects, but imposing the condition that the hedonic equilibrium is constant over time (except for inflation shifting it up). Thus, the 1995 and 2000 coefficients on TRI facilities are identical. As discussed by Kuminoff and Pope (2014), this model likely suffers from conflating movement along a hedonic price function with shifts in the function. The second row shows the same model, but allowing the coefficients to vary (Equation 9c). The 2000 coefficient increases slightly in absolute magnitude, so in this sample it appears conflation bias would lead to a downward bias. Results using alternative ½-mile or 2-mile grid cells are similar and are available upon request. The next two rows introduce time varying fixed effects, first by adding county-year fixed effects to the time-invariant 1-mile-grid effects, then by replacing them with year-by-2-mile-grid effects. The next four rows repeat these models, but using distance-weighted TRI sites within 2 miles rather than the raw count. All TRI coefficients are negative, as would be expected, and statistically significant.

The last column of Table 5 presents the lower bound welfare measure based on the 2000 coefficients, with bootstrapped confidence intervals. The bootstrap takes into account sampling error in estimation, forecasting (i.e. predicting direct effects), and reweighting tracts. The final column shows the estimated lower bound for the change in TRI sites in the LA area in billions of 2000 dollars. Excluding rows 1 and 5 (which impose constant coefficients), the values range from \$6.2b to \$8.6b, as a present value for what presumably is a permanent shock. With about 5.4 million households in the area covered, and at a discount rate of 5%, these values work out to about \$58-\$79 per LA household per year. Thus, in this case even the lower bound measure is substantial and may be informative for policy.

8. 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

willingness to pay 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 equivalent surplus. Simulations suggest the lower bound provides valuable information, on the order of 75% to 92% of the actual equivalent surplus.

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 Siegel 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 proposed here. In particular, 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 (2019). This would allow quasi-experimental methods to replace Rosen's two-stage strategy for non-marginal welfare measures.

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. 2019. Panel Data Hedonics: Rosen's First Stage as a "Sufficient Statistic." Working Paper.
- 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.
- Baum-Snow, N., and J. Marion. 2009. The Effects of Low Income Housing Tax Credit Developments on Neighborhoods. *Journal of Public Economics* 93: 654-66.
- 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.
- Dong, Y., and A. Lewbel. 2015. Identifying the effect of changing the policy threshold in regression discontinuity models. *Review of Economics and Statistics* 97(5): 1081-1092.
- Ekeland, I., J.J. Heckman, and L. Nesheim. 2004. Identification and Estimation of Hedonic Models. *Journal of Political Economy* 112(1): S60-S109.
- Figlio, D.N., and M.E. Lucas. 2004. What's in a Grade? School Report Cards and the Housing Market." *American Economic Review* 94: 591-604.
- Gamper-Rabindran, S., and C. Timmins. 2013. Does Hazardous Waste Remediation Raise Housing Values? *Journal of Environmental Economics and Management* 65(3): 345-60.
- Gopalakrishnan, S., M.D. Smith, J.M. Slott, and A.B. Murray. 2011. The value of Disappearing Beaches: A Hedonic Pricing Model with Endogenous Beach Width. *Journal of Environmental Economics and Management* 61(3): 297-310.
- 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*. Abingdon, UK: 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(4): 981-9.
- Keane, M.P., and K.I. Wolpin. 1997. The Career Decisions of Young Men. *Journal of Political Economy* 105(3): 473-522.
- Kennan, J., and J.R. Walker. 2011. The Effect of Expected Income on Individual Migration Decisions. *Econometrica* 79(1): 211-51.
- Klaiber, H.A., and V.K. Smith. 2013. Quasi Experiments, Hedonic Models, and Estimating Trade-offs for Local Amenities. *Land Economics* 89(3): 413-31.
- Kuminoff, N.V. 2012. Recovering Preferences from a Dual-Market Locational Equilibrium. Mimeo. <http://www.public.asu.edu/~nkuminof/K12.pdf>
- 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.
- Lind, R.C. 1973. Spatial Equilibrium, the Theory of Rents, and the Measurement of Benefits from Public Programs. *Quarterly Journal of Economics* 87(2): 188-207.
- 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.
- Palmquist, R.B. 2005. Weak Complementarity, Path Independence, and the Intuition of the Willig Condition. *Journal of Environmental Economics and Management* 49(1): 103-115.
- Parmeter, C.F., and J.C. Pope. 2013. Quasi-Experiments and Hedonic Property Value Methods. In *Handbook on Experimental Economics and the Environment*, ed. by J.A. List and M.K. Price. Cheltenham, UK: Edward Elgar: 3-66.
- Pope, Jaren C. 2008. Fear of Crime and Housing Prices: Household Reactions to Sex Offender Registries." *Journal of Urban Economics* 64(3): 601-14.
- Sieg, H., V.K. Smith, H.S. Banzhaf, and R. Walsh. 2004. Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods. *International Economic Review* 45(4): 1047-77.
- Smith, V.K., and H.S. Banzhaf. 2004. A Diagrammatic Exposition of Weak Complementarity and the Willig Condition. *American Journal of Agricultural Economics* 86(2): 455-66.
- Starrett, D.A. 1981. Land Value Capitalization in Local Public Finance. *Journal of Political Economy* 89(2): 306-327.
- Tchetgen Tchetgen, E.J., and T.J. VanderWeele. 2012. On Causal Inference in the Presence of Interference. *Statistical Methods in Medical Research* 21: 55-75.
- von Haefen, R.H. 2007. Empirical Strategies for Incorporating Weak Complementarity into Consumer Demand Models. *Journal of Environmental Economics and Management* 54: 15-31.
- VanderWeele, T.J., and E.J. Tchetgen Tchetgen. 2011. Effect Partitioning under Interference in Two-Stage Randomized Vaccine Trials. *Statistics and Probability Letters* 81: 861-9.
- Waugh, F.V. 1929. *Quality as a Determinant of Vegetable Prices*. New York: Columbia University Press.
- Yatchew, Adonis. 1998. Nonparametric Regression Techniques in Economics. *Journal of Economic Literature* 36: 669-721.

Figure 1. Defining Capitalization Effects

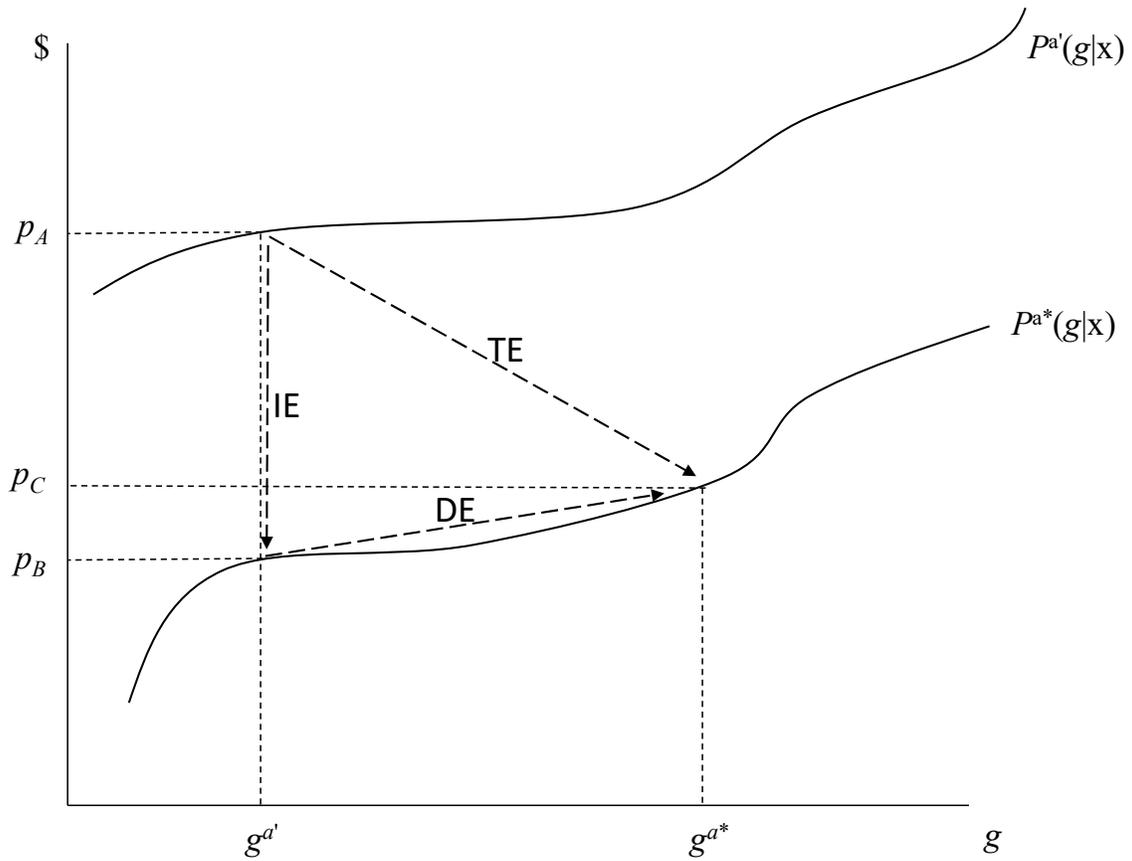


Figure 1 illustrates the indirect effect (IE), direct effect (DE), and total effect (TE) of a policy shifting the distribution of an amenity g . The policy shifts the hedonic price function from $p^a'(\cdot)$ to $p^a^*(\cdot)$. Even the price of untreated houses are affected by this shift, moving from 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

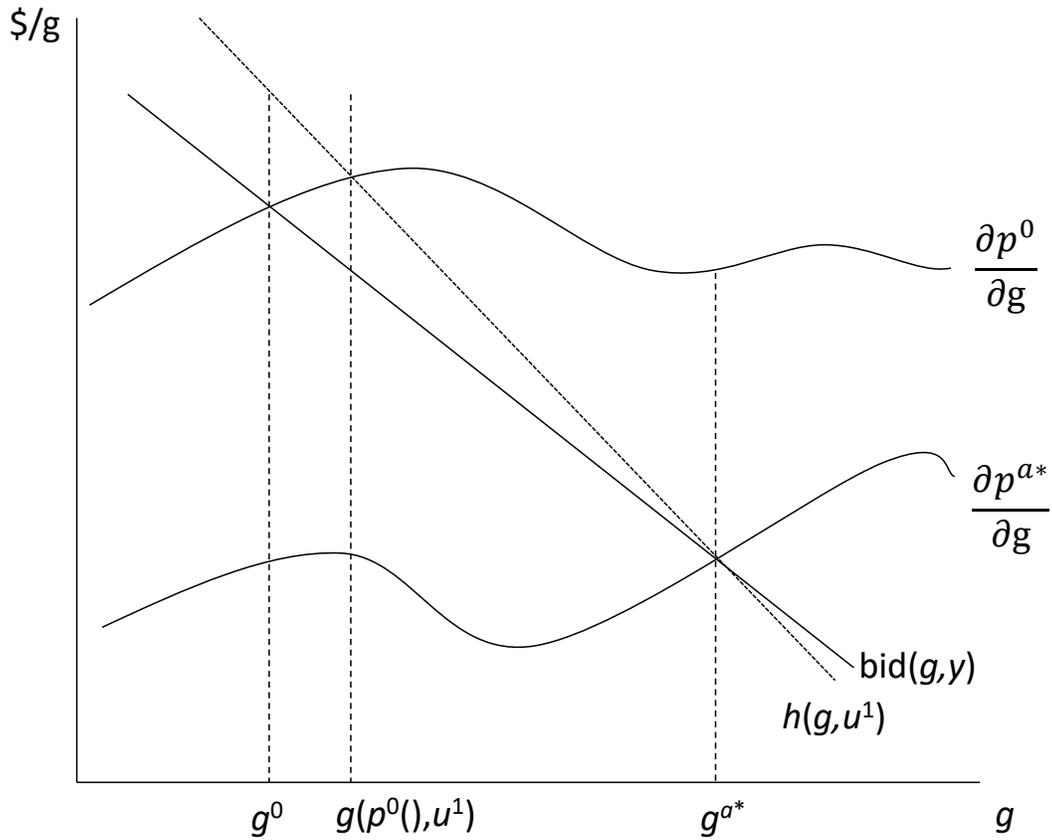


Figure 2 shows the Hicksian Equivalent Surplus (ES) as the area under the Hicksian demand curve $h(\cdot)$ from g^0 to g^{a*} . A partial-equilibrium version of the lower bound is illustrated by the fact that this area exceeds the movement along the hedonic, or the area $\int_{g^0}^{g^{a*}} \frac{\partial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

Figure 3. Pre-existing Trends in Housing Prices.

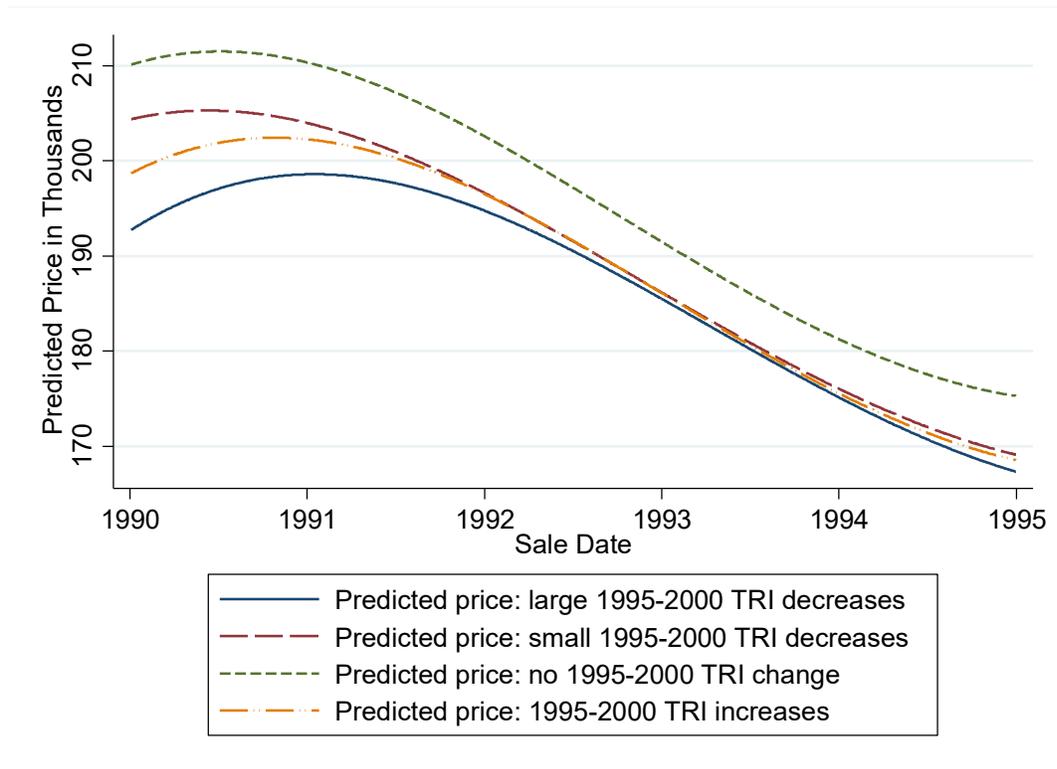


Figure 3 shows 1990-1995 mean predicted prices for four categories of houses, those with 1995-2000 decreases in TRI exposure of 2+ plants, a decrease of one plant, no change, and increases in exposure.

Table 1. 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 5th and 95th 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 2. 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 - 1,143)	\$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 - 1,354)	\$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 5th and 95th 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)

Table 3. 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 4. Pre-Existing Trends in Housing Prices

	Number of Facilities within 1 mile (1)	Weighted Number of Facilities within 1 mile (2)
Time Trend x 1995-2000 Change in facilities	-0.000123 (0.0000465)	-0.000272 (0.000150)
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 3, and 1-mile grid-cell fixed effects. Standard errors clustered by grid cell reported in parentheses.

Table 5. Results of Application to TRI data.

Measure of TRI Exposure	Model	1995 Coef. (Standard Error)	2000 Coef. (Standard Error)	Estimated Lower Bound, \$b (95% CI from bootstrap)
No. TRI facilities within 2 miles	1 mi grid FE (Time invariant Coef)	-0.01248 (0.0013)	-0.01248 (0.0013)	6.36 (5.40 – 7.32)
No. TRI facilities within 2 miles	1 mi grid FE	-0.00968 (0.0019)	-0.01492 (0.0015)	7.53 (6.18 – 8.19)
No. TRI facilities within 2 miles	1 mi grid FE + County-Year FE	-0.00866 (0.0020)	-0.01609 (0.0016)	8.22 (6.73 – 9.71)
No. TRI facilities within 2 miles	2 mi grid – Year FE	-0.00264 (0.0010)	-0.01606 (0.0015)	8.55 (7.33 – 9.78)
Distance-weighted	1 mi grid FE (Time invariant Coef)	-0.03399 (0.0035)	-0.03399 (0.0035)	5.37 (4.53 – 6.20)
Distance-weighted	1 mi grid FE	-0.02638 (0.0053)	-0.03976 (0.0042)	6.24 (5.12 – 7.37)
Distance-weighted	1 mi grid FE + County-Year FE	-0.02242 (0.0056)	-0.04353 (0.0045)	6.90 (5.63 – 8.17)
Distance-weighted	2 mi grid – Year FE	-0.01246 (0.0027)	-0.04281 (0.0040)	7.10 (6.03 – 8.16)

Each row corresponds to a separate regression using a different measure of TRI exposure and/or fixed effects. Standard errors clustered by grid cells. The final column shows the estimated lower bound on ES, in millions of dollars, for the LA area, from the observed 1995-2000 changes in TRI exposure.

APPENDIX. PROOF OF PROPOSITIONS.

Proof of Proposition 1.

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

$$(10) \quad dW = \sum_i [\tilde{e}_i(p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_h [(p^{a*}(\mathbf{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a'}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) - (c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}))].$$

The right side of Equation (10) can be decomposed as follows:

$$(17) \quad dW = \sum_i [\tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_i [\tilde{e}_i(p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - \tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*})] \\ + \sum_h (p^{a*}(\mathbf{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*})) + \sum_h (p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) \\ + \sum_h (p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}) - p^{a'}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) - \sum_h (c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})).$$

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

$$(18) \quad dW = \sum_i [\tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - e_i(p^{a*}(\cdot), u_i^{a*})] \\ + \sum_h [(p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) - (c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*}) - c(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'}))] \\ + \sum_h (p^{a*}(\mathbf{g}_h^{a*}, \mathbf{x}_h^{a*}) - p^{a*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a*})) \\ + \sum_i [(\tilde{e}_i(p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*}) - \tilde{e}_i(p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}), \mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}, u_i^{a*})) \\ - (p^{a'}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}) - p^{a*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}))].$$

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 (\mathbf{g}, \mathbf{x}) is held fixed is equal to the change in the price of the bundle (\mathbf{g}, \mathbf{x}) . Thus, the expression simplifies to:

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

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,

$$(19) \quad \sum_h (p^{a^*}(\mathbf{g}_h^{a^*}, \mathbf{x}_h^{a^*}) - p^{a^*}(\mathbf{g}_h^{a'}, \mathbf{x}_h^{a'})) \leq dW.$$

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

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:

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

It is similar to Equation (10), but now $u_i^{TC(R), a^*}$ is the utility achieved in scenario a^* when the household has to pay transaction costs from the reference point R . Additionally, the expenditure function $e_i^{TC(R)}(p^{a^*}(\cdot), u_i^{TC(R), a^*})$ is the same as in Equation (10) but it takes into account the expenditure necessary to pay the transaction cost. Likewise the constrained expenditure function $\tilde{e}_i^{TC(a^*)}(\cdot)$ takes such expenditures into account.

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

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

which parallels Equation (11).

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

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

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 $(\mathbf{g}_h^{a^*}, \mathbf{x}_h^{a^*})$ 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 $(\mathbf{g}_{i(a')}^{a'}, \mathbf{x}_{i(a')}^{a'})$ given the same price function and same reference point for transactions.

The second fact we will use is that

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

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:

$$\begin{aligned}
& \tilde{e}_i^{TC(a^*)} \left(p^{a^*} \left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right), \left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right), u_i^{TC(R),a^*} \right) \\
&= \tilde{e}_i^{TC(R)} \left(p^{a^*} \left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right), \left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right), u_i^{TC(R),a^*} \right) + (TC(R, a^*) + TC(a^*, a')) - TC(R, a') \\
&\geq \tilde{e}_i^{TC(R)} \left(p^{a^*} \left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right), \left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right), u \left(\left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right), y - p^{a^*} \left(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'} \right) - TC(R, a') \right) \right) \\
&\quad + [(TC(R, a^*) + TC(a^*, a')) - TC(R, a')] \\
&= e_i^{TC(R)} \left(p^{a^*}(\cdot), u_i^{TC(R),a^*} \right) + [(TC(R, a^*) + TC(a^*, a')) - TC(R, a')] \\
&\geq e_i^{TC(R)} \left(p^{a^*}(\cdot), u_i^{TC(R),a^*} \right).
\end{aligned}$$

The first equality follows from re-arranging Equation (20). Next, the inequality follows from Expression (22), plus the fact that the expenditure function is increasing in u . The next equality follows from the fact that $\tilde{e}_i^{TC(R)}(\cdot) = e_i^{TC(R)}(p^{a^*}(\cdot), u_i^{TC(R),a^*}) = y$. 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 $(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'})$ and numeraire consumption $y - p^{a^*}(\mathbf{g}_{i(a')}, \mathbf{x}_{i(a')}^{a'}) - TC(R, a')$, 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.