

Best Practices in Using Hedonic Property Value Models for Welfare Measurement

Kelly C. Bishop, Nicolai V. Kuminoff, H. Spencer Banzhaf, Kevin J. Boyle,
Kathrine von Gravenitz, Jaren C. Pope, V. Kerry Smith, and Christopher D. Timmins¹

Introduction

The hedonic property value model is among the most direct illustrations of how private markets can reveal consumers' willingness to pay for measures of environmental quality. There have been thousands of applications since the model was formalized in the 1970s and, if anything, the pace has accelerated due to advances in data accessibility, econometrics, and computing power. The hedonic model's enduring popularity is easy to understand. It seems like common sense by beginning with an intuitive premise that is both economically plausible and empirically tractable. The model envisions buyers choosing properties based on housing attributes (e.g., indoor space, bedrooms, bathrooms) and on location-specific amenities (e.g., air quality, park proximity, education, flood risk). In the absence of market frictions, spatial variation in amenities can be expected to be capitalized into housing prices. When buyers face the resulting menu of price-attribute-amenity pairings in the housing market, their purchase decisions can reveal their willingness to pay for marginal changes in each of the amenities.²

In recent years, the prevailing style of empirical hedonic research has evolved to incorporate insights from the "credibility revolution" in applied micro-econometrics. This revolution has raised expectations for data quality and econometric transparency. Recent research has refined our understanding of how parameters identified by quasi-experimental research designs map into welfare measures. This article distills the collective evidence from recent advances in the hedonic

¹ Bishop and Kuminoff share lead authorship. Bishop: Arizona State University, Department of Economics (Kelly.Bishop@asu.edu). Kuminoff: Arizona State University, Department of Economics and NBER (kuminoff@asu.edu). Banzhaf: Georgia State University, Department of Economics and NBER (hsbanzhaf@gsu.edu). Boyle: Virginia Tech, Department of Agricultural and Applied Economics and Virginia Tech Program in Real Estate (kjboyle@vt.edu). Von Gravenitz: ZEW – Leibniz Center for European Economic Research (Kathrine.vonGraevenitz@zew.de). Pope: Brigham Young University (Jaren_pope@byu.edu). Smith: Arizona State University and NBER (kerry.smith@cavecreekinstitute.com). Timmins: Duke University, Department of Economics and NBER (timmins@econ.duke.edu). We are grateful for helpful comments and suggestions from Noelwah Netusil and audiences at the World Congress of Environmental Economics.

² Nearly 100 years ago, in some of the first work on hedonic modeling, Waugh (1929, p.100) proposed a remarkably similar "statistical analysis". He suggested that "...instead of compiling the reported likes and dislikes of individuals, this type of statistical analysis attempts to estimate these preferences for the whole group of dealers or consumers in the market area by measuring the market price differentials due to a number of quality factors."

property-value literature to summarize current “best practices” for credible research designs and valid welfare interpretations. While hedonic property-value models are used for many purposes, our focus is on welfare measures that can be used to inform public policy.

We outline best practices when the researcher’s goal is to measure households’ willingness to pay (WTP) for a change in a spatially-varying amenity. These best practices start with a research design that identifies a clear source of exogenous variation in an amenity that prospective buyers can be assumed to observe. Microdata on the observed sale prices and the physical attributes of individual houses, together with location-specific measures for amenities are then used to estimate a housing-price function. With a flexible specification for the price function and conditions that assure the model’s assumptions are satisfied, the derivative of the price function will also be a function that can be interpreted as a measure for the amenity’s implicit price. The first-order conditions for utility maximization provide the conceptual basis for linking these implicit-price estimates to measures of household marginal WTP (MWTP). In principle, this process is straightforward. In practice, there are several important modeling decisions that must be made in order to define measures for sale prices and amenities and to select an econometric specification. Data limitations can also complicate identification and welfare interpretation. While the number of issues that must be considered in developing a “best practices” study may seem daunting, the effort is justifiable. The modern hedonic property-value model has been refined through more than forty years of intense scrutiny to become one of the premier approaches to valuing changes in environmental amenities in academic research, litigation, and public policy (Palmquist and Smith 2002, US EPA 2010).

The next section reviews the hedonic property-value model’s foundations. We organize our remaining discussion around issues that need to be addressed in estimating MWTP and in using these estimates to inform policy. We discuss how MWTP measures can be used to assess small changes in amenities and to bound welfare effects of large changes. While we note that MWTP measures can also be combined with additional information to estimate amenity demand curves, we leave the task of defining best practices in “second stage” hedonic demand estimation to future research.³ We conclude with a summary of what we feel are high-priority opportunities to advance the literature.

³ Kuminoff, Smith and Timmins (2013) review methods for using housing data to estimate welfare effects of large changes.

Most of the studies that helped to establish modern best practices in hedonic modeling used rich data on metropolitan housing markets in advanced economies. Data describing housing transactions, characteristics, and amenities are becoming increasingly available around the world, creating new opportunities to use hedonic models for policy analysis. An online appendix to this article summarizes data availability, data sources, and sample applications for 24 countries, and discusses additional modeling issues that may arise in rural areas and less-than-ideal data settings.

Conceptual Background

The hedonic framework has a long history in economics (Waugh (1929), Court (1939), Griliches (1961), and Lancaster (1971)). Rosen's (1974) seminal paper established the hedonic framework as an equilibrium model for understanding what differentiated products' prices could reveal to market observers. It offered a direct analog to the Cowles Commission's logic for connecting structural and reduced-form models (Morgan (1990)).⁴ In the housing context, the primitives of the hedonic model consist of the supply for housing, including developers' decisions for new home construction and factors influencing re-sales of existing homes, and the joint distribution of household preferences and income. After buyers and sellers negotiate transactions, market equilibrium occurs when no agent can increase utility by moving. This equilibrium concept implies a relationship between house prices and characteristics that reveals each buyer's MWTP for each house characteristic under the assumptions that buyers are fully informed, freely mobile, and able to purchase continuous levels of each characteristic.⁵ We use four graphs in Figure 1 to explain key features of the model.⁶

Panel 1a plots housing price as a function of the measure for one of the local amenities, e.g., an environmental amenity, holding the physical characteristics and other location-specific features constant.⁷ Panel 1b illustrates the process through which the model will allow the price function to reveal buyers' MWTP. This panel adds two buyers' bid curves. These bid curves each trace out the maximum amount that each buyer is willing to pay as a function of the amenity level (holding

⁴ It is also closely connected to Tiebout's (1956) seminal analysis of local public goods provision.

⁵ These assumptions have subtle implications. For example, "free mobility" does not imply that it must be costless to move. Buyers' choices will still reveal their MWTP at the time of their purchase decisions if they were able to choose continuous levels of each characteristic facing a fixed cost of moving. Meanwhile, the assumptions may be violated if some prospective buyers are excluded from renting or buying properties because of discrimination.

⁶ Readers seeking a technical exposition should see Palmquist (2005), Kuminoff, Smith and Timmins (2013), Freeman, Herriges and Kling (2014), Phaneuf and Requate (2017) or Taylor (2017).

⁷ In Rosen's model, each point on the price function is the tangency between a particular seller's offer curve and a particular buyer's bid curve. These are the points at which market trades occur. We suppress sellers' offer curves in Figure 1b in order to focus on demand.

other influences on their choices constant). Purchases occur where bid curves are tangent to the price function. Buyer 2 purchases a house with amenity level A_2 at a price of P_2 and buyer 1 purchases a less expensive house in a lower-amenity area. These two coordinate sets, (P_1, A_1) and (P_2, A_2) , are the points at which each buyer's MWTP for a small change in the amenity (bid-curve slope) equals the amenity's implicit price (price-function slope). Panel 1c illustrates this same result in a different way: by explicitly distinguishing what is revealed by a hedonic equilibrium. First, notice that the algebraic form used for the hedonic price function implies an algebraic form for the implicit price function for the amenity, since this is implied by the derivative of the hedonic price function. The dashed demand curves are not observed by the analyst. With this in mind, Panel 1c illustrates that home purchases only reveal information about demand at the points of intersection where demand curves for the amenity intersect the implicit price function for the amenity A (P^A). Assuming this structure effectively describes the underlying market mechanics, one can recover each buyer's MWTP for the amenity in three steps: (1) use observed sales data to estimate the hedonic price function, (2) partially differentiate the price function with respect to the amenity of interest to recover the implicit price function, and then (3) evaluate the implicit price function for each buyer to recover their MWTP.

Panel 1c also illustrates another important point: the implicit price is just that – a price. Just as prices in a "regular market" are equal to MWTP at a point on the demand curve, so too the implicit price derived in step (2) reveals only a marginal value, not the entire demand function. Notice that any number of flatter or steeper demand specifications could be drawn through the point (P_1^A, A_1) . Following the logic of the Cowles' approach to demand estimation, applied by Rosen to the hedonic model, additional information would be needed to infer amenity demand from the implicit price function or, equivalently, to predict welfare effects of counterfactual, non-marginal changes in amenity levels. Researchers have developed several strategies for providing this additional information, which we will return to in a later section. For now, we focus on using the implicit price function to recover MWTP.

Finally, panel 1d illustrates another important point. It shows how demand curves and implicit price functions may change over time in response to changes in the distribution of amenity levels and market primitives. For example, at some point after the initial period S a new policy may lead to changes in the amenity levels throughout a market. Meanwhile, some households may experience wealth shocks. These types of changes can induce migration and alter the market-clearing

price function and its gradient (the implicit price function). The year-S implicit price function identifies MWTP for the distribution of buyers in year S and the same is true for buyers in year T, one year after the policy change. However, the year-S and year-T distributions of MWTP will generally differ. In this example, we assume that buyer 1 has left the market so we do not include a year-T demand function in the graph. Buyers who can be observed multiple times, such as buyer 2, may have experienced changes in their personal circumstances (demand shocks). Understanding temporal changes in the price function and the distribution of homebuyers' amenity demand curves can be important for econometric estimation of implicit amenity prices and for mapping implicit prices into measures of MWTP associated with changes in amenities.

Best Practices for Estimating Marginal Willingness to Pay

Market Definition

The model's conceptual logic implies that the market should be chosen to satisfy the "law of one price function". That is, when a house can be fully defined by a unique bundle of physical characteristics and location-specific amenities, then equivalent bundles sell for the same price throughout that market. The precise spatial and temporal boundaries that satisfy these conditions may vary across geography and over time as information, institutions, and moving costs change. A common approach is to define the market as a single metropolitan area over a few years (e.g. Pope 2008b, Abbott and Klaiber 2013). An alternative is to pool data over larger areas and longer periods, while modeling the parameters of the hedonic price function as evolving over space and time (e.g. Kuminoff and Pope 2014, Walls et al. 2017).

In principle, moving costs could lead to violations of the law of one price function. However, for households that move within metropolitan areas, moving costs are unlikely to vary substantially with the destination locations. This is due to the fact that the physical and financial moving costs (e.g., realtor fees and truck rentals) do not vary with within-metropolitan-area destination locations and the psychological moving costs are limited by the fact that within-metropolitan-area moves typically allow households to maintain ties to family, friends, and neighborhoods. Consequently, the "law of one price" can be maintained by arbitrage between locations within a metropolitan area. In contrast, moving between metropolitan areas may impose much larger moving costs. Equally important, workers who move between metropolitan areas may be forced to change jobs.

Potential variation across geographic areas in tax policy and the cost of living (aside from housing) adds further complications. Because the hedonic property-value model abstracts from labor-market considerations and heterogeneous moving costs, focusing on larger geographic areas which include multiple metropolitan areas can undermine the conceptual logic conventionally used for mapping a hedonic price function into MWTP measures. Researchers can avoid this problem by using data on commuting patterns to determine when moving to a different metropolitan area would likely imply moving to a new job.

Pooling data over a long period such as a decade or more introduces similar problems, as there is unlikely to be arbitrage over time. Housing-price functions can evolve during boom-bust cycles as macroeconomic factors change the amounts homebuyers are willing to pay for amenities. Homebuyers' MWTP may also evolve with changes in information and policy. For instance, a policy that improves air quality may reduce homebuyers' MWTP for further improvements (i.e., by moving them down their demand function for air quality) whereas a policy that improves homebuyers' knowledge about the negative health effects of pollution may increase their MWTP for clean air (i.e., by shifting out their demand function). Including time-specific dummy variables as "intercept shifters" in the hedonic price function may help to control for housing price inflation but will be insufficient to control for temporal changes that contribute to shifts in MWTP, as the hedonic equilibrium evolves.

In principle, some sources of spatio-temporal variation in the shape of hedonic price functions can be addressed through econometric flexibility (e.g., McMillen and Thorsnes (2003)). However, parametric assumptions used to model geographic and temporal changes dictate how the results should be interpreted. Researchers can relax parametric assumptions when pooling data over multiple metropolitan areas and years by using interactions between time dummies, geography dummies, and price function parameters to allow price functions to differ across space and time. We discuss some of the issues related to econometric specification in the following sections.

Overall, narrowing the assumed extent of the market will tend to improve internal validity by increasing the likelihood that the "law of one price function" holds, but it may reduce external validity and the ability to study geographically coarse amenities, such as climate features. If the goal is to understand how amenities affect residential sorting across metropolitan areas (where a different form of arbitrage, not the simplest version of the law of one price best describes the equilibrium process) then the class of Tiebout sorting models summarized by Kuminoff, Smith,

and Timmins (2013) provides one means for consistently incorporating job opportunities and moving costs.

Data Collection

The gold standard in data collection is to obtain a random sample (or the universe) of housing-transaction prices and characteristics for the relevant study area.⁸ Most studies focus exclusively on single-family houses. In most parts of the United States, for example, housing transactions are a matter of public record and are usually filed with county tax-assessment boards.⁹ This type of access enables researchers to work with data that approximate the universe of single-family housing sales in specific time periods.

Some studies use data on sales of undeveloped land or data on rental rates for houses and apartments. Estimating a hedonic model of vacant-land sales is consistent with the idea that the price function maps how prices vary with land characteristics. However, there are important institutional factors, such as zoning, prior easements, and access to public water supplies that affect how land may be used. Housing rents may be used in conjunction with sales prices by converting sales prices into annualized user-costs of housing using standard formulas in the literature (Poterba (1984) and Himmelberg, Mayer, and Sinai (2005)). However, rental rates can present added complications relative to data on single-family transactions in that there may be ambiguity about important rental-contract features, such as which party pays for utilities and maintenance. The short-term nature of rental contracts may also weaken the incentive for renters to become fully-informed about local amenities prior to entering the market. On the other hand, rentals may better reflect current amenity flows, allowing the analyst to avoid the dynamic problem that buyers face, as buyers are potentially forward-looking over future amenity flows. The use of rental-rate data may also be particularly important for recovering unbiased measures of average MWTP in neighborhoods where rates of owner occupation may be low.

In recent years, data on housing transaction prices, characteristics, and amenities have become increasingly available for large portions of other countries, including Australia, Japan, and South Korea. A few countries, such as Denmark and Sweden, have additionally granted researchers access to administrative records containing rich socioeconomic panel data on buyers and sellers. In

⁸ We note the potential selection bias that may occur when focusing only on houses that sell. Gatzlaff and Haurin (1998) propose a correction procedure that uses information on the non-price characteristics of houses that do not sell.

⁹ Cleaned data that had been previously filed with county boards can be purchased from CoreLogic, ATTOM Data Solutions, and other vendors.

other countries, such as Canada and Portugal, it is still difficult to obtain microdata on transactions. Appendix Table 1 (online supplementary material) provides a country-by-country summary of what we could determine about data availability, data sources, and sample applications for Australia, Austria, Belgium, Canada, Chile, China, Denmark, Finland, France, Germany, Hungary, Ireland, Italy, Japan, Netherlands, Norway, Poland, Portugal, South Korea, Spain, Sweden, Switzerland, United Kingdom, and the United States. While this set of countries is far from comprehensive, it provides a starting point for researchers looking for housing data or sample applications. The appendix also discusses additional challenges that may arise in less than ideal data settings including regulation of prices, sparse transactions, and lack of transaction prices.

Data preparation

It is reasonable to expect that publicly-available data on housing sales, often collected for other reasons, will include some data-entry errors as well as some sales which do not arise in a competitive bidding process. Identifying these cases and dropping them reduces the scope for measurement error. For instance, it is common to exclude transactions in which the buyer and seller share the same last name and, therefore, have a higher probability of being related. It is also common to drop foreclosure sales and purchases by real-estate investment firms, as there is a higher probability that the property has characteristics or quality-issues not documented in transactions data. Finally, it is common to remove outliers that embed seemingly-conspicuous data entry mistakes (e.g., a house with 1,800 bedrooms or 3 square feet). Many researchers address these outliers by dropping a small fraction of sales with the highest and lowest values for each characteristic. Since there is no commonly-accepted threshold for what defines an outlier, it is important to document these types of decisions and assess the sensitivity of findings to them.

Predicted prices

Some studies rely upon prices that have been predicted in some form or another. One example is that Census data sets often include a self-reported “value”. These data are generated from survey questions that ask occupants how much they think their properties would sell for if they were to sell it. Another example is predicted prices from property assessors and other companies (e.g., Zillow). Transaction prices are always preferable to predicted prices. The problem with predicted

prices is that they embed measurement error which varies systematically over time, buyer demographics, housing characteristics, and neighborhood amenities (Banzhaf and Farooque 2013).¹⁰ This correlation may lead to bias in price-function-parameter estimates.

Spatially-aggregated data

Individual records of transaction prices are preferable to spatially-aggregated measures, such as mean or median prices within Census tracts, zip codes, or counties. A fundamental drawback of using such measures is that theory does not establish an equilibrium mapping from price aggregates to aggregate welfare measures. Regressing mean prices on mean amenity levels may yield an unbiased estimate of the hedonic price function, but not of the implicit price function (i.e., the derivative) which is necessary to recover MWTP.¹¹ Moreover, median prices are not necessarily equal to the price level at the median attributes, so hedonic price regressions using medians may be biased from the outset. Hence, using summary statistics as if they were transaction level prices can undermine welfare measurement.

Even apart from problems of interpretation, median prices have been found to introduce measurement error that can bias price-function parameters. For example, in an application to Superfund cleanups of hazardous-waste sites, Gamper-Rabindran and Timmins (2013) find that focusing on the median house within a neighborhood recovers lower estimates for MWTP than does using individual house records. This result occurs because hazardous-waste sites tend to be located near the lowest-price houses within a neighborhood and, as a result, the benefits of cleanups are mainly capitalized into housing prices at the lowest quantiles of the housing-price distribution. Problems with medians are not limited to studying hazardous-waste sites or other point-source amenities. Banzhaf and Farooque (2013) use data for Los Angeles to compare several commonly-used measures for house prices and find that among all the commonly-used measures, medians have the weakest correlation with local public goods, income, and other indices. While the reasons for this finding are not fully understood, it suggests that the results for hedonic models based on median prices should be interpreted with caution.

¹⁰ The problem may be especially pronounced when prices are predicted using an algorithm. In this case a hedonic regression may simply recover a re-construction of the algorithm.

¹¹ The derivative of the price function evaluated at the mean amenity level does not equal the mean of the derivatives, e.g., let $P_i^A = a + bA_i$ equal the implicit price function where $i = [1,2]$. In this case, $\bar{P}^A = \frac{P_1^A + P_2^A}{2}$ and $\bar{A} = \frac{A_1 + A_2}{2}$ and this result is easily verified.

Assigning amenity levels to houses

The model's revealed-preference logic requires the analyst to characterize how buyers perceive the amenity levels at each residential location. This task can present multiple challenges. One challenge is to develop an objective measure of spatial variation in the amenity that can be matched to individual houses. This process can be complicated by the often "patchy" nature of amenity data. For instance, air-quality monitoring networks generate data on pollution levels, measured as ambient concentrations defined as a statistic, e.g., the mean over a specific time period at a point in space. Since houses are naturally situated in gaps between monitoring stations, analysts must use spatial interpolation, air-dispersion models, or predictions from satellites to assign pollution levels to houses. Likewise, proximity to recreation sites such as beaches, lakes, and parks may be measured by geographical distance, by driving distance, by total travel time, or by the share of land devoted to that recreation use within some geographic area around a house. Analysts must decide which measure best reflects the landscape characteristics that matter to homebuyers.

Another challenge is to consider whether homebuyers' subjective beliefs about the amenity's dispersion coincide with objective measures and, if not, to consider alternative ways of modeling buyer beliefs. The broader, nonmarket-valuation literature suggests that subjective beliefs about environmental quality do not always coincide with objective measures (Boyd et al. (2015)). This concern also appears to hold for private attributes of many market goods outside of the environmental literature. Thus, it can be important to document the information channels that may influence buyers' beliefs and assess sensitivity to the choice among candidate measures of the amenity. For instance, Davis (2004) demonstrates that the measurement of cancer risk associated with a cluster of leukemia cases is robust to different assumptions about how homebuyers formed beliefs about the evolving level of risk, including basing it on the cumulative number of cancer cases, the cumulative number of newspaper articles about the cancer cluster, or a Bayesian rule for risk updating. In contrast, Pope (2008) found that a new law requiring real estate agents to disclose information about airport noise caused housing prices to adjust around an international airport. His findings suggest that the disclosure rule changed prospective buyers' beliefs about the spatial dispersion of noise. Under the hypothesis that the information disclosure improved buyers' knowledge of noise levels, sales after the post-disclosure period would be expected to provide more accurate estimates for MWTP.

The dynamic process that leads to changes in amenities presents another challenge. Because housing is a durable asset, households' purchase decisions may also reflect their expectations about the future evolution of local-amenity levels. When buying a house in the current period, for example, a forward-looking household would compare both the current and anticipated future flows of amenities to the current purchase price. In cases where amenity levels exhibit a mean-reversion (mean-diversion) in their trend and households are responding to it, the current implicit price of the amenity would underestimate (overestimate) households' MWTP. The conventional approach to this problem is to take a moving average of the amenity level through time. More recently, Bishop and Murphy (2018) show how to map hedonic-price-function estimates into households' MWTP when buyers are forward-looking over changing amenity levels. They clarify when a dynamic approach would be warranted.

Econometric Specification

Functional form

Theoretical and simulation evidence suggest that the hedonic price function should be assumed to be nonlinear. For example, Ekeland, Heckman, and Nesheim (2004) demonstrate that transactions between heterogeneous buyers and sellers in a differentiated-product market yield an equilibrium hedonic price-function gradient that is generically non-linear in characteristics.¹² Monte Carlo simulations by Cropper, Deck, and McConnell (1988) show that relatively-flexible specifications for the price function, such as Box-Cox transformations, provide more accurate estimates of average MWTP than simpler linear and log-linear specifications in realistic data environments when there are no omitted house characteristics. Moreover, Kuminoff, Parmeter, and Pope (2010) found that flexible Box-Cox models continue to outperform linear models in the presence of omitted variables when spatial dummies are used to mitigate omitted-variable bias. Semiparametric and nonparametric methods can provide additional flexibility in estimating hedonic price functions and their gradients though there is currently no simulation evidence on the implications of the bias-variance tradeoff implicit in such approaches (e.g., Bajari and Benkard (2005), Parmeter, Henderson, and Kumbhakar (2007), McMillen and Redfearn (2010), Bishop and Timmins (2018)).¹³ Yet another

¹² In other words, it would take extraordinarily strong assumptions about the shapes of utility functions and housing-production functions to generate linear hedonic price functions as equilibrium outcomes.

¹³ Simulation methods could be used to investigate bias-variance tradeoffs between parametric nonlinear models, semiparametric models, and nonparametric models to guide functional form decisions in future research.

reason for using nonlinear functional forms is to allow the description used for the market equilibrium to reflect complementarity between amenities. For example, the implicit price of proximity to a public park may vary with the levels of crime, noise, and air quality (Albouy, Christensen, and Sarmiento-Barbieri (2018)).

Econometric errors

As a rule, applications both rely on robust, “sandwich” variance-covariance estimators for the standard errors of the hedonic-price-function parameters and cluster at a spatiotemporal scale consistent with the variation in the amenity of interest. This practice is motivated by the common belief that it is prudent to assume the errors may exhibit heteroskedasticity and autocorrelation (spatial and temporal). Some applications have also experimented with spatial-weighting models that are analogous to feasible generalized least squares. While such models may enhance small-sample efficiency, provided the true parametric form of the error-correlation structure (e.g., nearest-neighbor weighting versus distance decay) is known, without this information, standard errors may be biased. As a rule, most applications avoid making such assumptions and rely on large sample sizes to justify invoking the asymptotic properties of robust, clustered, estimates for the standard errors.

Mitigating Omitted-Variable Bias and Related Threats to Identification

Theory and empirical evidence suggest that amenities will be spatially correlated due to natural features of geography, environmental feedback effects, and voting on local public goods. This potential for correlation fuels widespread concern about omitted-variable bias. There are two aspects of this concern: first, it seems reasonable to assume that analysts can never include every amenity that matters to buyers and, second, latent locational characteristics are likely to be correlated with the amenity of interest, causing bias. For instance, if wealthy and well-educated homebuyers move to areas with better air quality and then vote to increase public-school funding, estimates of MWTP for air quality will be biased upward if school quality is omitted from the model. Thus, a research design that isolates exogenous variation in the amenity of interest becomes necessary to ensure credibility of the resulting estimates. The literature has implemented many such approaches. We separately discuss many of the commonly-implemented research designs in this section.

These approaches amend the theoretical logic underlying the conventional hedonic framework in ways that introduce important tradeoffs in research designs. On the one hand, it is important to assure econometric credibility by mitigating threats due to omitted-variable problems. On the other hand, the economic foundations for the conventional hedonic framework provide a clear basis for recovering policy-relevant welfare measures. Sometimes efforts to enhance econometric credibility direct attention to quasi-experimental designs and/or data on predicted prices that compromise the ability to interpret the estimates as measures of MWTP. In these cases, the analyst has a defensible estimate of an effect but cannot link that estimate to a meaningful economic concept, limiting its policy relevance. We discuss this tradeoff with each research design.

Difference-in-Difference Research Designs

Numerous studies have leveraged environmental-policy changes and natural experiments as quasi-random shocks to amenities and used this interpretation to identify how the shocks lead to changes in housing prices. This process is colloquially described as “capitalization,” although the word’s precise meaning has evolved over time. Econometric models of the capitalization process generally fit within a difference-in-difference (DID) framework. This group includes fixed-effect and first-difference estimators that utilize repeated transactions for the same houses. It also includes estimators that pool repeated cross-sections of transactions from the same geographic market. Sometimes these models utilize instrumental variables and/or regression-discontinuity designs. Overall, the DID framework is distinguished by the way it analyzes transactions before, relative to after, a change to the spatial distribution of the amenity.

The DID framework’s main strength is to mitigate omitted-variable bias by isolating quasi-random variation in amenities over time. Chay and Greenstone’s seminal 2005 study demonstrates how non-attainment of the U.S. Environmental Protection Agency’s standards on maximum-allowable particulate-matter concentrations at the county level can be used as an instrument in analyzing how spatially-varying reductions in particulate matter between 1970 and 1980 influenced the median, self-reported, property value at the county level. Thus, their study is notable for the innovation in research design but is subject to concerns regarding the use of median prices, the use of predicted prices, and the assumption that the law of one price function holds across the United States from 1970 to 1980. In addition to mitigating concerns about omitted amenities, their instru-

mental-variables approach addresses potential bias from measurement errors in air-pollution levels. The quasi-experimental logic from this study has since been adapted to micro data and transaction prices to estimate capitalization effects of changes in a wide range of amenities including cancer risk (Davis 2004), fracking externalities (Muehlenbachs et al. (2015)), air pollution reductions (Bento, Freedman, and Lang (2015)), sand dune construction (Dundas (2017)) and open space (Lang (2018)) to name only a few.

The most common approach to DID is to assume a stationary price function. Figure 1d illustrates the main challenge in interpreting results from this approach: the price functions change over time. Indeed, theory suggests that environmental policies and other events that create quasi-experimental changes in amenity levels will also cause price functions to adjust. Thus, the prices of houses in the “control group” are affected by the policy, as well as prices of the “treated” houses. This result is distinct from any changes that may occur due to macroeconomic forces and other background events during the study period. The key issue is not whether price functions change, but rather whether the changes that occur are small enough to ignore. Large changes in price functions that are not specifically modeled can undermine welfare analysis.

Kuminoff and Pope (2014) show that when a price function shifts over time, the standard DID model ignoring the shift will yield biased estimates of the slopes of the price function and thus biased estimates of MWTP. The standard model mixes information from two equilibria into one estimate, but the economic logic of the hedonic model involves arbitrage at a point in time, not across time. Thus, the model conflates shifts in the price function with movements along the price function. Mixing these two effects becomes problematic if the change in amenities is correlated with the change in price from the shift. In an application to public school quality, Kuminoff and Pope show that price functions in five U.S. metropolitan areas changed during a 5-year boom period from 2003 to 2007. Importantly, this was the same time that the “No Child Left Behind Act,” a policy which targeted school quality and sought to improve information conveyed to parents, was implemented. Incorrectly assuming a time-constant price function would have produced a 75% downward bias in MWTP estimates, in part because school-quality changes were correlated with baseline school-quality levels. This outcome contrasts with a 94% upward bias in cross-section models that ignore the omitted variable problem. In a simulation-based study of housing market equilibria that compared capitalization to MWTP, Klaiber and Smith (2013) find that naïve

cross-section models may or may not outperform more sophisticated DID research designs if the latter are compromised by price-function changes.

One solution is to impose additional exclusion restrictions and to instrument for the change in the amenity of interest. As Kuminoff and Pope (2014) discuss, this strategy has the potential to identify MWTP in the post-shock period. However, it depends on strong orthogonality conditions: the instruments must be orthogonal to baseline conditions (amenities and structural characteristics), as well as to changes in unobservables. A second strategy is to model the change in the price function, generalizing the DID model by interacting price function parameters with time-period dummies. Kuminoff, Parmeter and Pope (2010) provide simulation evidence that this strategy improves the accuracy of estimates for MWTP in both the pre-shock and post-shock periods, and it has been implemented in recent empirical studies (e.g., von Gravenitz 2018). Banzhaf (2018) further shows that this approach can identify a lower bound on a general-equilibrium welfare measure under much weaker econometric assumptions. The results hold for open cities and in the presence of moving costs. A third approach is to assume a parametric form for the dynamic evolution of omitted variables and assume that buyers have rational expectations for that process (Bajari et al. 2011).

Matching Estimators

Matching estimators seek to mitigate omitted-variable bias by matching houses that received an amenity treatment with a set of untreated control houses. The goal is to find control houses that are as similar as possible to each treated house in terms of observed and unobserved physical and locational characteristics. This process essentially uses observed property characteristics to control for unobserved characteristics in a more flexible way than what would result from a simple linear OLS regression. The challenge is to determine the precise criteria for selecting matches. While the econometric properties of matching estimators have been thoroughly analyzed in the program-evaluation literature, the accuracy of their estimates for MWTP for amenities in housing market data has yet to be evaluated in the simulation frameworks that provide insights on other methods.

A notable feature of matching applications is that amenity treatment is typically discrete, such as whether a house provides lake access (Abbott and Klaiber (2013)) or whether it has received an Energy-Star certification (Walls et al. (2017)). Discrete treatment violates the hedonic model's assumption that buyers can choose continuous levels of each amenity, conditional on all other

variables. While this raises a concern that the tangency condition underlying welfare interpretation is violated, the simulation studies cited above generally find that such violations are not necessarily fatal in the sense that the margin of error in estimates for average MWTP for integer measures of housing attributes (e.g., number of bedrooms or bathrooms) is often similar to the margin of error for continuous attributes (e.g., square footage or proximity to parks).

Including spatial dummy variables

Transactions data often offer large sample sizes so it is possible to include a set of spatial dummy variables for local neighborhoods such as school districts, zip codes, or Census tracts in the U.S. Kuminoff, Parmeter, and Pope (2010) show that this strategy can mitigate omitted-variable bias and improve the accuracy of estimates for MWTP in both linear and non-linear specifications of the hedonic price function. Selecting the geographic scale for the dummies presents a tradeoff: shrinking neighborhood size strengthens the case for absorbing omitted variables (reducing bias) but reduces the identifying variation in the amenity (increasing variance). GIS maps showing how the amenity varies within and across candidate neighborhoods can help to diagnose this tradeoff (e.g., von Gravenitz (2018)). Abbott and Klaiber (2011) show how the Hausman-Taylor (1981) estimator may be used to judge the importance of this bias-variance tradeoff.

Boundary-Discontinuity Research Designs

Boundary-discontinuity designs seek to sharpen the spatial dummy-variable approach by leveraging variation in amenity levels within a neighborhood. The idea is to identify an amenity's marginal implicit price from sharp changes in amenity levels that occur at administrative or geographic boundaries. By limiting the estimation sample to houses located within close proximity to boundaries (e.g., within a quarter mile) and using dummy variables for neighborhoods around each boundary to absorb all of the omitted amenities common to both sides, this strategy assumes a sharp difference in the amenity of interest is most likely to lead to a price differential. Sample applications include school attendance zone boundaries (Black (1999)), flood zone boundaries (Pope (2008b)), and public water-service-area boundaries (Muehlenbachs, Spiller, and Timmins (2015)).

While the boundary-discontinuity design is consistent with the hedonic model's conceptual underpinnings, it presents at least two challenges. First, household sorting may confound the identification strategy by generating sharp differences across boundary zones in latent endogenous amenities. For instance, if wealthier households tend to locate on the "high quality" side of a school-zone boundary and also tend to divert more money to neighborhood parks, this would lead to a biased estimate of MWTP for school quality if parks are not included in the model. Bayer, Ferreira, and McMillan (2007) found that when they address this problem by controlling for differences in socioeconomic status of households on opposite sides of boundaries there is a first order change in their estimates of MWTP for school quality. The second challenge is that the resulting measures of MWTP may be limited in their ability to inform policies affecting the broader housing market; the analyst must be willing to assume that the identified part of the housing price surface is representative for the population of interest. While it is common to assume that the boundary neighborhoods are representative of the broader population, we are unaware of any evidence that evaluates these assumptions.

Assessing Robustness

Every study embeds modeling decisions that affect welfare conclusions. We have discussed the choice of the amenity variable, the source of variation in the amenity, the decisions made about sample composition (including observations removed as likely coding errors and outliers), and the parametric assumptions stemming from selecting a specification for the price function. Reporting the sensitivity of welfare conclusions to these and other modeling decisions can help to mitigate concerns that results are driven by arbitrary assumptions or by outlying observations. Dundas (2017, Figure 5) provides an informative graphical example of robustness within a targeted sensitivity analysis.

Best Practices for Using MWTP Estimates to Inform Policy

Incorporating Heterogeneity

Figures 1a and 1b illustrate how differentiating the hedonic price function with respect to the amenity of interest yields an implicit price function for the amenity that may be evaluated for each household's observed selection. These values provide point estimates for each household's MWTP. If it is possible to match housing-transaction records to administrative data on households,

then the heterogeneity in MWTP estimates may be linked to buyers' demographic characteristics (e.g., race, income, education, children). In the United States, this has been done by merging publicly-available, Home-Mortgage-Disclosure-Act data describing basic demographic characteristics of buyers who finance their purchases with federally-insured mortgages with data describing housing transactions (e.g., Bishop and Timmins (2017)). In some European countries, researchers can link property transactions with even richer government records on household demographics (von Graevenitz (2018)).

Demand Estimation for Non-Marginal Changes in the Amenity

As noted earlier, Figure 1c illustrates that a buyer's single observed consumption choice reveals only one point on the buyer's amenity-demand curve. Thus, the demand function, which would be required to measure welfare for non-marginal changes, cannot be recovered for each household. Rosen (1974) proposed estimating the demand function by regressing MWTP estimates on corresponding quantities consumed and demand shifters such as buyers' demographic characteristics. Unfortunately, this estimation strategy presents an endogeneity problem (see Bartik 1987 and Epple 1987), as unobserved tastes simultaneously determine both a buyer's MWTP (at the point of consumption) and their chosen quantity.

Thus, to recover the structure of demand, the literature has developed a variety of econometric strategies. These include imposing restrictions on: the parametric form of utility (e.g., Bajari and Benkard (2005)); the scope of preference heterogeneity (e.g., Ekeland, Heckman, and Nesheim (2004), Bishop and Timmins (2019)); the stability of preference heterogeneity across cities (Bartik (1987) and Zabel and Kiel (2000)); the stability of household preferences over time (Bishop and Timmins (2018) and Banzhaf (2017)); and the extent of spatial sorting (Bartik (1987) and Zhang, Boyle, and Kuminoff (2015)). These studies recover the demand function in a variety of ways. Several provide proof-of-concept applications. Yet, no individual approach has been clearly adopted by the empirical literature as a best-practices approach for amenity-demand estimation.

Non-Parametric Bounds and Approximations for Willingness to Pay

Some studies use hedonic price function estimates for back-of-the-envelope calculations that multiply MWTP by large changes in amenities. For these calculations to be exact, demand must be

perfectly elastic and invariant to the changes. The likelihood that such assumptions provide reasonable approximations decreases with the size of the change. When the change in the amenity is believed to be too large to assume a perfectly elastic demand curves, an alternative approach to estimation of the demand function is to obtain non-parametric bounds and approximations following the logic of Varian (1982). Consider Figure 1d. Suppose that a household's demand does not change between periods. In period T, household 2 is observed to pay an implicit price of P_3^A to consume A_3 units of the amenity. In Period S, this household pays an implicit price of P_4^A to consume A_4 units. If, for example, we had only observed the period-T choice, the rectangle defined by $P_3^A \times (A_4 - A_3)$ would provide an upper bound for the willingness to pay that would have been recovered by integrating under any downward sloping demand curve going through that point and over to A_4 . Naturally, the lower bound would be zero. If instead we had only observed the period-S choice, the rectangle defined by $P_4^A \times (A_4 - A_3)$ would provide a lower bound for the willingness to pay to avoid a policy that would decrease household 2's amenity level to A_3 from A_4 . Naturally, the upper bound would be infinity. Without additional functional-form restrictions, this is all that one can say about household-specific welfare from a single data point. Note that these bounds correspond to indifference curves that range between the cases where money and the amenity are perfect substitutes to perfect complements.

When the household is observed at two or more points, these bounds can be made tighter (Varian 1982). Moreover, with additional functional form restrictions, the demand curve can be calculated without any econometric estimation. As Bajari and Benkard (2005) and Bishop and Timmins (2018) point out, if we assume the demand curve is linear then two points are enough to identify it without any statistical estimation. One can simply connect the dots between (P_3^A, A_3) and (P_4^A, A_4) . By extension, three points are enough to identify a quadratic demand curve, and so forth. Furthermore, Banzhaf (2017) shows that, with two points, the connect-the-dots approach provides a second-order approximation to a true Hicksian welfare measure, regardless of the unknown Hicksian demand curve. Moreover, even if one cannot follow households over time, with a panel of houses one can either identify this second-order approximation with additional structure or alternatively bound it. Bishop and Timmins (2018) illustrate the connect-the-dots approach for air quality in the Bay Area of California, finding considerable heterogeneity in WTP that would be missed with the conventional approach.

Opportunities to Advance the Literature

Employing Administrative Records

Increased access to administrative records offers the potential to improve our understanding of buyers' revealed preferences for local amenities. First, enhanced information about buyers, including their demographic characteristics, working situation, income levels, and wealth, can enable researchers to analyze how hedonic estimates for MWTP vary with these factors (e.g., von Graevenitz 2018). Second, as some structural models of residential-sorting behavior rely on strong distributional assumptions (Kuminoff, Smith, and Timmins (2013)), greater access to administrative records would enhance our ability to recover the distributional implications of policies and would allow distributional assumptions to be evaluated. Finally, the information contained in administrative records may help in identifying demand curves for an amenity. For instance, consider household 2 in the lower right panel of Figure 1d that is observed in both period S and period T under different implicit price schedules; in principle, one should be able to literally "connect the dots" and recover this household's demand curve, as long as the household's income and preferences remain constant (Bishop and Timmins (2018)). Administrative records could assist in isolating the households who "fit" this assumption. Similarly, the data sets used by Voorheis (2017) and Bishop, Ketcham, and Kuminoff (2018) to track the long-term evolution of individuals' health and wealth could, in principle, be matched to housing transactions to yield new insights on how changes in health and wealth affect the demand for specific amenities.

Focusing on External Validation for MWTP Function Estimation

Another way to advance the empirical literature would be to determine which of the existing approaches to demand estimation is most suitable for policy analysis by testing the assumptions used to move from point estimates of MWTP to a demand function. For example, Galiani, Murphy, and Pantano (2015) illustrate how modeling assumptions may be tested in a discrete-choice model by developing testable, out-of-sample predictions about how changes in housing prices induce households to adjust their neighborhoods. In principle, this methodology could be adapted to propose testable predictions for how changes in local amenities would induce households to adjust their neighborhoods.

Adapting the Policy Elasticity to Hedonic Models

Recently, public economists have argued that Hendren's (2015) concept of a marginal value of public funds (MVPF) offers the means for economists "...to harness the fruits of the 'credibility revolution' for the public finance goal of welfare analysis" (Finkelstein (2018), p.1). While the MVPF concept is not new to the field of environmental economics, adapting it to the hedonic model has potential to enhance the relevance of MWTP estimates for assessing the efficacy of competing environmental policies. In the environmental context, the MVPF is defined as the willingness to pay for a marginal change in an amenity relative to the net incremental cost of providing that change through a policy. This metric differs from a simple benefit-cost ratio by incorporating the causal impact of behavioral responses to the policy on the government budget. For example, a policy that reduces air pollution may also reduce federal health-care expenditures or raise property tax revenues (by increasing property values). The challenge for future research would be to credibly estimate sufficient statistics for describing how these behavioral responses affect taxpayers (i.e., the policy elasticities) and then combine this information with hedonic estimates for MWTP and data on policy implementation costs.¹⁴ If this could be done for multiple policies, the resulting MVPF measures could be used to order policies according to their return on investment and determine the most efficient way to allocate a marginal increase in government expenditures on the environment. The hedonic equilibrium framework in Banzhaf (2017, 2018) could provide a starting point for further research on the relationship between MWTP and MVPF.

Investigating Heterogeneity in Beliefs

Finally, there is evidence that consumers often have heterogeneous beliefs about product attributes, with some consumers being misinformed at the time of purchase even when it comes to high-stakes financial decisions such as choosing a college major (Wiswall and Zafar (2015)), choosing a health insurance plan (Handel and Kolstad (2015)), or developing a strategy to save money for retirement (Bernheim, Fradkin, and Popov (2015)). The same is true for expensive durable goods such as cars (Busse et al. (2015)), refrigerators (Houde (2017)), and water heaters (Allcott and Sweeney (2017)). When consumers are not fully-informed about product characteristics, their choices can fail to reveal their preferences. Several of these studies overcome the problem by incorporating survey data on consumers' beliefs. Adapting such approaches to the hedonic property-

¹⁴ Chetty (2009) and Heckman (2010) discuss sufficient statistics in other contexts.

value context has potential to improve the accuracy of welfare measures for changes in amenities if buyers are not fully-informed.

Evidence on the degree to which buyers are informed about the characteristics of their houses and neighborhoods is mixed. Myers (2017) uses capitalization effects of fuel-price changes to conclude that homebuyers are well-informed about how future energy costs vary with a house's heating technology (gas versus oil). In contrast, Pope (2008a, 2008b) uses capitalization effects of real-estate information disclosures to conclude that some homebuyers did not pay attention to publicly-available information about flood risk and airport noise prior to mandatory disclosure laws that required them to sign forms stating their awareness of the amenities. Bakkensen and Barrage (2017) provide more direct evidence by surveying homebuyers about their beliefs of their flood risk. Their findings suggest residents of more flood-prone areas are more likely to underestimate flood risk. Since the current evidence suggests that the accuracy of buyers' beliefs varies from context to context, future research on housing purchases should consider adapting the methods that have been developed to incorporate heterogenous beliefs. The challenge is to improve our understanding of households' information sets and to use this information to refine welfare measures. A potential first step would be to combine transactions data with surveys that reveal buyers' beliefs about the spatial dispersion of amenities at a point in time. The next step would be to learn how households' beliefs evolve over time as they process new information. Ma (2018) demonstrates how one can model the learning process within a structural dynamic sorting model, finding that doing so has a large impact on estimates for the willingness to pay for brownfield remediation.

Conclusion

The choice of where to live is perhaps the single most important decision affecting a household's consumption of environmental amenities. It is natural to expect that housing markets would have the ability to reveal information about the welfare effects of changes in amenities. The hedonic property-value model provides an economically-plausible and empirically-tractable way to distil this information. The model's credibility and policy relevance have been improved in recent years by advances in applied theory and econometrics. While we noted many cautions regarding implementation, these cautions should not deter readers from using the model. The literature has simply generated more knowledge about the "do's" and "don'ts" of hedonic modeling than for most other microeconomic frameworks used to analyze policy. The model's success is underscored by the

fact that it is used in the real world to inform public and private decisions.¹⁵ The bottom line is that modern “best practices” property-value regressions can provide credible estimates for what households are willing to pay for environmental amenities.

References

- Abbott, Joshua K. and H. Allen Klaiber. 2011. “An Embarrassment of Riches: Confronting Omitted Variable Bias and Multiscale Capitalization in Hedonic Price Models”. *Review of Economics and Statistics*, 93 (4): 1331-1342.
- Abbott, Joshua K. and H. Allen Klaiber. 2013. “The Value of Water as an Urban Club Good: A Matching Approach to Community-Provided Lakes”. *Journal of Environmental Economics and Management*, 65: 208-224.
- Albouy, David, Peter Christensen and Ignacio Sarmiento-Barbieri. 2018. “Unlocking Amenities: Estimating Public-Good Complementarity”. *NBER Working Paper #25107*.
- Allcott, Hunt and Richard L. Sweeney. 2017. “The Role of Sales Agents in Information Disclosure: Evidence from a Field Experiment.” *Management Science*, 63(1): 21-39.
- Bajari, Patrick and C. Lanier Benkard. 2005. “Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach.” *Journal of Political Economy*. 113(6): 1239-1276.
- Bajari, Patrick, Jane Cooley Fruehwirth, Kyoo Il Kim and Christopher Timmins. 2012. “A Rational Expectations Approach to Hedonic Price Regressions with Time-Varying Unobserved Product Attributes: the Price of Pollution”. *American Economic Review*, 102(5): 1898-1926.
- Bakkensen, Laura and Lint Barrage. 2017. “Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics: Going Under Water?” *NBER Working Paper #23854*.
- Banzhaf, H. Spencer. 2017. "Panel Data Hedonics: Rosen's First Stage as a 'Sufficient Statistic'." Working Paper.
- Banzhaf, H. Spencer. 2018. "Difference-in-Differences Hedonics." Working Paper.

¹⁵ Triplett (1990) reported the first use of hedonic methods in a government price index in 1968 for new one-family houses sold. More recently (2007) he summarized the use of hedonic methods in developing prices indexes and other measures for the national accounts. In the private sector there are many areas where hedonic pricing is used in a proprietary fashion, most obviously in real estate by companies such as Zillow.

- Banzhaf, H. Spencer, and Omar Farooque. 2013. "Interjurisdictional Housing Prices and Spatial Amenities: Which Measures of Housing Prices Reflect Local Public Goods?" *Journal of Regional Science and Urban Economics*. 43: 635-648.
- Bayer, Patrick, Fernando Ferreira and Robert McMillan. 2007. "A Unified Framework for Measuring Preferences for Schools and Neighborhoods." *Journal of Political Economy*. 115(4): 588-638.
- Bartik, Timothy J. 1987. "The Estimation of Demand Parameters in Hedonic Price Models." *Journal of Political Economy*, 95(1): 81-88.
- Bento, Antonio, Matthew Freedman, and Corey Lang. 2015. "Who benefits from environmental regulation? Evidence from the Clean Air Act Amendments." *Review of Economics and Statistics* 97(3): 610-22.
- Bernheim, B. Douglas, Andrey Fradkin, and Igor Popov. 2015. "The Welfare Economics of Default Options in 401(k) Plans." *American Economic Review*, 105(9): 2798-2837.
- Bishop, Kelly, Jonathan Ketcham, and Nicolai Kuminoff. 2018. "Hazed and Confused: The Effect of Air Pollution on Dementia." *NBER Working Paper #24970*.
- Bishop, Kelly and Alvin Murphy. 2019. "Incorporating Dynamic Behavior into the Hedonic Model." *Review of Economics and Statistics*, 101(1): 134-145.
- Bishop, Kelly and Christopher Timmins. 2019. "Estimating the Marginal Willingness to Pay Function Without Instrumental Variables." *Journal of Urban Economics*, 109(2019): 66-83.
- Bishop, Kelly and Christopher Timmins. 2018. "Using Panel Data to Easily Estimate Hedonic Demand Functions." *Journal of the Association of Environmental and Resource Economists*, 5(3): 517-543.
- Black, Sandra E. 1999. "Do Better Schools Matter? Parental Valuation of Elementary Education." *Quarterly Journal of Economics*, 114(2): 577-599.
- Boyd, James, Paul Ringold, Alan Krupnick, Robert J. Johnston, Matthew A. Weber and Kim Hall. 2015. "Ecosystem Services Indicators: Improving the Linkage between Biophysical and Economic Analyses." *RFF Discussion Paper #15-40*.
- Busse, Meghan R., Devin G. Pope, Jaren C. Pope, and Jorge Silva-Risso. 2015. "The Psychological Effect of Weather on Car Purchases." *Quarterly Journal of Economics*. 130(1): 371-414.

- Chay, Kenneth Y. and Michael Greenstone. 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy*. 113(2): 376-424.
- Chetty, Raj, 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics* 1, 451–88.
- Cropper, Maureen L., Deck, Leyland B., McConnell, Kenneth E., 1988. "On the choice of functional form for hedonic price functions." *Review of Economics and Statistics*. 70 (4), 668–675.
- Court, Andrew T. 1939. "Hedonic Price Indexes with Automotive Examples." *The Dynamics of Automobile Demand*, General Motors Corporation, NY, NY. 99-117.
- Davis, Lucas. 2004. "The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster." *American Economic Review*, 94(5): 1693-1704.
- Dundas, Steve. 2017. "Benefits and Ancillary Costs of Natural Infrastructure: Evidence from the New Jersey Coast." *Journal of Environmental Economics and Management*, 85: 62-80.
- Ekeland, Ivar, James J. Heckman, and Lars Nesheim. 2004. "Identification and Estimation of Hedonic Models." *Journal of Political Economy*, 112(1): S60-S109.
- Epple, Dennis. 1987. "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products." *Journal of Political Economy*, 95(1): 59-80.
- Finkelstein, Amy. 2018. "Welfare Analysis meets Causal Inference: A suggested Interpretation of Hendren (2016, 2017)" Unpublished lecture notes, March.
- Freeman III, A. Myrick, Joseph A. Herriges and Catherine L. Kling. 2014. *The Measurement of Environmental and Resource Values: Theory and Methods*. RFF Press.
- Galiani Sebastian, Alvin Murphy, and Juan Pantano. 2015. "Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment." *American Economic Review* 105(11):3385–3415.
- Gamper-Rabindran, Shanti, and Christopher Timmins. 2013. "Does Cleanup of Hazardous Waste Sites Raise Housing Values: Evidence of Spatially Localized Benefits." *Journal of Environmental Economics and Management*. 65: 345-360.
- Gatzlaff, Dean H. and Donald R. Haurin. 1998. "Sample Selection and Biases in Local House Value Indices." *Journal of Urban Economics*. 43(2): 199-222.
- von Graevenitz, Kathrine. 2018. "The Amenity Cost of Traffic Noise". *Journal of Environmental Economics and Management*, 90: 1-22.

- Griliches, Zvi. 1961. "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change." *The Price Statistics of the Federal Government: Review, Appraisal, and Recommendations*, ed. Price Statistics Review Committee, NBER, NY, NY. 173–96.
- Handel, Benjamin R. and Jonathan T. Kolstad. 2015. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review*, 105(8): 2449-2500.
- Hausman, J.A., and W.E., Taylor, 1981. "Panel Data and Unobservable Individual Effects." *Econometrica*, 49, 1377-1398.
- Heckman, James J. 2010. "Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy." *Journal of Economic Literature*, 48:356-98.
- Hendren, Nathaniel. 2015. "The Policy Elasticity" NBER working paper #19177.
- Himmelberg, C., C. Mayer, and T. Sinai. 2005. "Assessing High House Prices: Bubbles, Fundamentals, and Misperceptions." *Journal of Economic Perspectives*. 19(4): 67–92.
- Klaiber, H. Allen and V. Kerry Smith. 2013. "Quasi-Experiments, Hedonic Models, and Estimating Tradeoffs for Local Amenities." *Land Economics*. 89: 413-431.
- Kuminoff, Nicolai V. Christopher F. Parmeter, and Jaren C. Pope. 2010. "Which Hedonic Models Can We Trust to Recover the Marginal Willingness to Pay for Environmental Amenities?" *Journal of Environmental Economics and Management* 60 (3), 145-160.
- Kuminoff, Nicolai V. and Jaren C. Pope. 2014. "Do Capitalization Effects Measure the Willingness to Pay for Public Goods?." *International Economic Review*, 55(4), 1227-1250.
- Kuminoff, Nicolai V., V. Kerry Smith, and Christopher Timmins. 2013. "The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets." *Journal of Economic Literature* 51(4): 1007-1062.
- Houde, Sébastien. 2017. "How Consumers Respond to Product Certification and the Value of Energy Information." *Rand Journal of Economics*, Forthcoming.
- Lancaster, Kelvin. 1971. "Consumer Demand: A New Approach." Columbia University Press, NY, NY.
- Lang, Corey. 2018. "Assessing the Efficiency of Open Space Provision." *Journal of Public Economics*, 158: 12-24.
- Ma. Lala. 2018. "Learning in a Hedonic Framework: Valuing Brownfield Remediation". *Forthcoming in International Economic Review*.

- McMillen, Daniel P., and Christian L. Redfearn. 2010. "Estimation and Hypothesis Testing for Nonparametric hedonic house price functions." *Journal of Regional Science* 50(3): 712-733.
- McMillen, Daniel P., and Paul Thorsnes. 2003. "The aroma of Tacoma: time-varying average derivatives and the effect of a superfund site on house prices." *Journal of Business and Economic Statistics* 21(2): 237-46.
- Morgan, Mary S. 1990. *The History of Econometric Ideas*. Cambridge University Press.
- Muehlenbachs, Lucija, Elisheba Spiller, and Christopher Timmins. 2015. "The Housing Market Impacts of Shale Gas Development". *American Economic Review*, 105(12): 3633-3659.
- Myers, Erica. 2017. "Are Home Buyers Myopic? Evidence from Capitalization of Energy Costs". *American Economic Journal: Economic Policy*, Forthcoming.
- Palmquist, Raymond B. and V. Kerry Smith. 2002. "The Use of Hedonic Property Value Techniques for Policy and Litigation." *The International Yearbook of Environmental and Resource Economics*, 2003, 115-64.
- Palmquist, Raymond B. 2005. "Property Value Models." in *Handbook of Environmental Economics, Volume 2*. edited by Karl Gören-Mäler and Jeffery Vincent. North Holland Press.
- Parmeter, Christopher F., Daniel J. Henderson, and Subal C. Kumbhakar. 2007. "Non-parametric Estimation of a Hedonic Price Function." *Journal of Applied Econometrics*, 22(3), 695-699.
- Phaneuf, Daniel J. and Till Requate. 2017. *A Course in Environmental Economics: Theory, Policy and Practice*. Cambridge University Press.
- Pope, Jaren C. 2008a. "Buyer Information and the Hedonic: The Impact of a Seller Disclosure on the Implicit Price for Airport Noise." *Journal of Urban Economics*. 63: 498-516.
- Pope, Jaren C. 2008b. "Do seller Disclosures Affect Property Values? Buyer Information and the Hedonic Model." *Land Economics*. 84(4): 551-572.
- Poterba, James M. 1984. "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach." *Quarterly Journal of Economics*. 4: 729-752.
- Rosen, Sherwin. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy*, 82(1): 34-55.
- Taylor, Laura. 2017. "The Hedonic Method," in *A Primer on Nonmarket Valuation*. Edited by Patricia A. Champ, Kevin Boyle and Tomas C. Brown. Springer Netherlands.

- Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy*, 64(5): 416-24.
- Triplett, Jack E. 1990. "Hedonic Methods in Statistical Agency Environments: An Intellectual Biopsy," in *Fifty Years of Economic Measurement*. Edited by Ernst R. Berndt and Jack E. Triplett. Chicago: University of Chicago Press for NBER, Studies in Income and Wealth, Vol 54.
- Triplett, Jack E. 2007. "Zvi Griliches's Contributions to Economic Measurement" Edited by Ernst R. Berndt and Charles R. Hulten. *Hard to Measure Goods and Services: Essays in Honor of Zvi Griliches*. Chicago: University of Chicago Press for NBER, Studies in Income and Wealth, Vol 67.
- United States Environmental Protection Agency. 2010. Guidelines for Preparing Economic Analyses. EPA-240-R-10-001.
- Varian, Hal R. 1982. "The nonparametric approach to demand analysis." *Econometrica*:50(4): 945-973.
- Voorheis, John. 2017. "Air Quality, Human Capital Formation and the Long-term Effects of Environmental Inequality at Birth." *Working paper*.
- Walls, Margaret, Todd Gerarden, Karen Palmer, and Xian Fang Bak. 2017. "Is Energy Efficiency Capitalized into Home Prices? Evidence from Three U.S. Cities." *Journal of Environmental Economics and Management*. 82: 104-124.
- Waugh, Frederick V. 1929. *Quality as a Determinant of Vegetable Prices: A Statistical Study of Quality Factors Influencing Vegetable Prices in the Boston Wholesale Market*, Ph.D thesis, Columbia University 1929.
- Wiswall, Matthew and Basit Zafar. 2015. "Determinants of College Major Choice: Identification using an Information Experiment". *Review of Economic Studies*. 82(2): 791-824.
- Zabel, Jeffrey E., and Katherine A. Kiel. 2000. "Estimating the demand for air quality in four US cities." *Land Economics* 76(2): 174-194.
- Zhang, Congwen, Kevin J. Boyle, and Nicolai V. Kuminoff. 2015. "Partial Identification of Amenity Demand Functions." *Journal of Environmental Economics and Management* 71(1), 180-197.

Fig 1a: The hedonic price function in amenity space




Fig 1b: The buyers' purchase decisions




Fig 1c: The implicit price function reveals buyers' MWTP

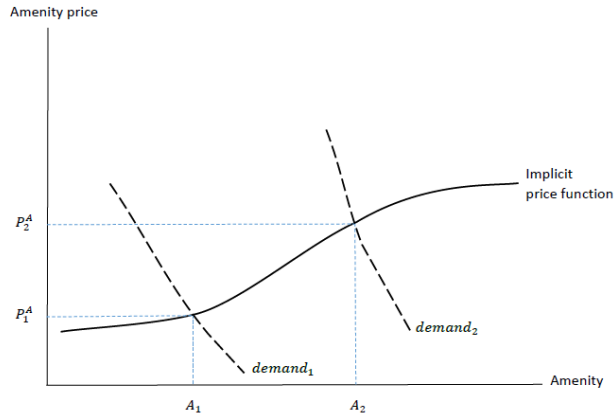


Fig 1d: The implicit price function and MWTP change over time




Figure 1: Using the Hedonic Price Function to Infer Buyers' MWTP for an Amenity

Note: Fig 1a shows the price of housing increasing in the level of an amenity. Fig 1b shows purchase decisions for two households. Household 2 purchases a house providing amenity A_2 at price P_2 . This is the point at which the household's demand curve intersects the amenity's implicit price function, as shown in Fig 1c. Fig 1d shows this relationship changing after a policy changes amenity levels. See the main text for additional detail.