

Rohlf, Chris; Sullivan, Ryan; Kniesner, Thomas J.

Working Paper

Hedonic Estimation under Very General Conditions Using Experimental and Quasi-Experimental Designs

IZA Discussion Papers, No. 7554

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Rohlf, Chris; Sullivan, Ryan; Kniesner, Thomas J. (2013) : Hedonic Estimation under Very General Conditions Using Experimental and Quasi-Experimental Designs, IZA Discussion Papers, No. 7554, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<http://hdl.handle.net/10419/89843>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 7554

Hedonic Estimation under Very General Conditions Using Experimental and Quasi-Experimental Designs

Chris Rohlfs
Ryan Sullivan
Thomas J. Kniesner

August 2013

Hedonic Estimation under Very General Conditions Using Experimental and Quasi-Experimental Designs

Chris Rohlfs

Morgan Stanley

Ryan Sullivan

Naval Postgraduate School

Thomas J. Kniesner

*Syracuse University,
Claremont Graduate University and IZA*

Discussion Paper No. 7554

August 2013

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Hedonic Estimation under Very General Conditions Using Experimental and Quasi-Experimental Designs¹

This paper develops a generalized hedonic model in which an exogenous shock to a single product attribute can affect other attributes, the markets for the product's complements and substitutes, and aggregate quantity produced. These factors are shown to be empirically relevant and to cause bias in traditional approaches. Experimental and quasi-experimental estimators of attribute demand are introduced that address these biases, are transparent, and are straightforward to implement. One of these estimators is applied to measure the marginal military recruit's valuation of educational benefits, which is found to range across packages from $-\$0.024$ to $+\$0.467$ per dollar of benefits.

JEL Classification: D12, C35, C31, D61, C9

Keywords: hedonic, discrete choice, identification, experiment, quasi-experiment, marginal willingness to pay, attribute demand, amenity, heterogeneous goods, generalized, endogenous attributes, complement, substitute

Corresponding author:

Ryan Sullivan
Defense Resources Management Institute
205 Halligan Hall
Naval Postgraduate School
Monterey, CA 93943
USA
E-mail: rssulliv@nps.edu

¹ Special thanks to Gus Bartuska and Qu Feng for expert research assistance; thanks also to Dan Black, Gregorio Caetano, Bill Horrace, Boyan Jovanovic, Tom Kniesner, Jeff Kubik, Derek Laing, Derek Neal, Jan Ondrich, Stuart Rosenthal, Kevin Tsui, Andy Vogel, Pete Wilcoxon, Paul Wilson, and seminar participants at Clemson, New York Camp Econometrics, NYU, Rochester, and Syracuse for many helpful comments.

Hedonic estimation, or the measurement of “marginal willingness to pay” (MWTP) for product attributes, is a vital tool for measuring the benefits of public policies that improve safety or environmental, school, or health care quality (*e.g.*, Black, 1999; Chay and Greenstone, 2005; Cutler, Rosen, and Vijan, 2006; Viscusi, 1993, 1996). Hedonic methods are used to understand the demand for heterogeneous goods such as automobiles, computers, food, housing, and jobs (*e.g.*, Bajari and Benkard, 2005; Hamermesh, 1999; Kiesel and Villas-Boas, 2007; Raff and Trajtenberg, 1995; Sheppard, 1999) and to calculate the Consumer Price Index and one fifth of expenditures in the Gross Domestic Product (Landefeld and Grimm, 2000; Moulton, 2001).

Reading the hedonics and discrete choice literature gives the impression that MWTP cannot be identified without imposing highly restrictive assumptions about utility and markets, even when a natural experiment is available. The existing models assume that unobserved product attributes are uncorrelated with observed ones or do not exist (Berry, Levinsohn, and Pakes, BLP, 1995; Epple, 1987; Rosen, 1974). In quasi-experimental settings, this assumption rules out “endogenous attribute” effects, through which an exogenous shock to a housing attribute such as local school quality affects neighborhood composition, crime, and the types of peers in that school. These models also assume that the product of interest has no complements or substitutes, so that a location-specific attribute like weather cannot affect the labor market as well as the housing market.² They also assume that the aggregate quantity consumed is exogenous and does not respond to price changes.³

There is a widespread belief in the literature that these restrictions are appropriate and necessary to estimate MWTP. Applied studies of heterogeneous goods generally employ slight modifications of these frameworks or measure reduced-form price effects without estimating MWTP. Much of the recent empirical work in hedonic estimation focuses on quasi-experiments, and some innovative studies have incorporated quasi-experimental variation into existing hedonic or discrete choice models (*e.g.*, Bayer, Ferreira, and McMillan, 2007; Berry and Haile, 2010; Chay and Greenstone, 2005; Klaiber and Smith, 2009; Kuminoff and Pope, 2010a, 2010b; Lewbel, 2000; Parmeter and Pope, forthcoming).⁴ No previous hedonic or discrete choice frameworks allow for unobserved product attributes that are affected by

² Roback (1982) allows for this one exception of housing and labor however, the literature generally ignores this and other possible interactions, such as effects of location-specific attributes on the prices of local goods.

³ Rosen (1974) and Epple (1987) additionally require that markets are “thick,” so that every conceivable product is available, and that supply is perfectly competitive. BLP (1995) additionally require specific functional forms for utility and firm costs and a specific distribution for heterogeneity in preferences.

⁴ These studies relax some important assumptions typically used in the literature but are generally faithful to those previous models. Additionally, some recent theoretical studies relax the functional form assumptions from these models but leave the frameworks largely intact (*e.g.*, Athey and Imbens, 2007; Ekeland, Heckman, and Nesheim, 2004; Heckman, Matzkin, and Nesheim, 2010).

exogenous shocks, complements or substitutes with the good of interest, and aggregate quantities that vary.⁵

The stylized frameworks used in previous hedonic and discrete choice studies are valuable as economic models. They shed light on the tradeoffs that occur in markets for heterogeneous goods and the patterns that one might expect to see in the data. For the purposes of measurement and policy evaluation, however, it is desirable to have robust hedonic estimators whose results are correct generally, and not only in the special cases described by these models. This paper demonstrates the identifiability of MWTP without imposing any of the restrictions described above, and it presents simple and straightforward estimators of MWTP and related measures for use in experimental and quasi-experimental settings.

Section II of this study provides an intuitive discussion of the types of bias that endogenous omitted attributes, complement and substitute goods, and aggregate quantity effects generate in traditional approaches. That section also presents evidence from the housing and labor markets in Alaska that these biases are empirically relevant.

Section III summarizes the estimators developed in this study to address these biases. Of the seven estimators presented in this study, the first ones use the least restrictive modeling assumptions but have the most demanding data requirements. The modeling assumptions are more restrictive and the data requirements less demanding with successive estimators, and the last one presented is the closest to existing quasi-experimental hedonic estimators.

Section IV adapts Jones's (1984) general model of the demand for heterogeneous goods to a hedonic framework. That section introduces concepts, notation, and a minimal set of assumptions necessary to define MWTP. The focus of this study is identification, and additional assumptions are not imposed to ensure the existence of equilibrium. The model is compatible with competitive or strategic equilibria provided that such additional assumptions are imposed.

Sections V and VI describe in detail the experiments and quasi-experiments that are proposed in this study. Section V develops nonparametric experimental estimators that identify the entire distribution across consumers of the demand for a given product attribute. Section VI develops experimental and quasi-experimental estimators of the aggregate demand for a product attribute among a population of consumers. The experimental estimators used in this study avoid the effects of endogenous omitted attributes and complement and substitute goods by offering products and subsidies to small numbers of consumers. The quasi-experimental estimators exploit variation across consumers or economies to disentangle the direct effects of a change in a product attribute on quantity demanded of a good and the

⁵ Only Roback (1982) allows for one type of complementarity (housing and jobs), and only Bartik (1987) considers the case in which aggregate quantities vary. BLP (1995) allow for market shares (but not aggregate quantity produced) to vary. No framework has addressed more than one of these biases.

indirect effects of the change on other attributes and prices, which have their own effects on quantity demanded of the good.

The estimators developed in this paper rely upon straightforward identification proofs, are transparent, and are considerably easier to implement than are previous hedonic and discrete choice estimators. Variations on these estimators have been applied in recent studies to estimate the value of freedom from jail, the demand for avoiding the Vietnam draft, the value of a statistical life, and the demand for class size reductions in elementary school (Abrams and Rohlfs, 2011; Rohlfs, 2012; Rohlfs, Sullivan, and Kniesner, 2013; Rohlfs and Zilora, 2013). Section VII applies one of the estimators of aggregate demand for an attribute to measure the value that military recruits place on funds for higher education. Contributions to an education account are found to have been valued at \$0.24 to \$0.47 per dollar of benefits, and direct tuition and stipend support is found to have been valued at -\$0.02 to \$0.08 per dollar of benefits. Section VIII concludes.

II. ILLUSTRATION OF BIAS IN HEDONIC MODELS

To illustrate the sources of bias that the current study addresses, let $Price_{ht}$ be the average price of house h in year t . Let z_{1ht} be an observable attribute about house h , such as local school quality. The value of z_{1ht} is determined by a quasi-experiment so that it varies exogenously across locations and over time. Let z_{2ht} be an attribute about house h that is difficult to measure, such as the pleasantness of neighbors in the area. $Price_{ht}$ is a linear function of the two attributes and an error term denoting unobserved attributes:

$$(1) \text{ } Price_{ht} = \beta_0 + \beta_1 z_{1ht} + \beta_2 z_{2ht} + \varepsilon_{ht}$$

The aim of a hedonic price regression in this case is to identify β_1 , the effect of attribute z_{1ht} on housing prices, holding all other attributes constant. Once identified, this price effect is used in a second-stage procedure to estimate MWTP for z_{1ht} (Epple, 1987; Rosen, 1974).⁶ For this second stage to produce accurate estimates, the estimates from the hedonic regression must be consistent.

The assignment of z_{1ht} is essentially random, and z_{1ht} is consequently uncorrelated with any predetermined factors that influence $Price_{ht}$. An improvement in school quality, however, will cause affluent and educated people to move into the area. Supposing that being affluent and educated is correlated with being a pleasant neighbor, the school quality improvement will directly affect the pleasantness of neighbors. If z_{2ht} is not included as a control in the regression, OLS estimates of Equation 1 measure the reduced-form effect of the shock to school quality on $Price_{ht}$ and not the structural parameter β_1 . This reduced-form effect includes the direct effect of z_{1ht} and the indirect effects

⁶ The procedure proposed by BLP (1995) is different from that described here, but BLP require consistent estimation of the effect of the attribute on the decision to purchase the product.

of z_{1ht} through the mechanism of z_{2ht} . Hence, OLS produces a biased estimate of β_1 , and the magnitude of the bias is $\beta_2 * cov(z_{1ht}, z_{2ht})$. Even if data were available on z_{2ht} , the values of this second attribute were not assigned through a quasi-experiment and are likely to be correlated with ε_{ht} . Consequently, without an instrument for z_{2ht} , an OLS regression of $Price_{ht}$ on z_{1ht} and z_{2ht} will produce biased estimates of β_2 , and by failing to adequately control for z_{2ht} , the regression will also produce biased estimates of z_{1ht} .⁷

The pleasantness of neighbors is not unique in this regard. Many location-specific attributes are influenced by the composition of local residents and businesses. Such attributes will vary in response to any exogenous shock that causes consumers or firms to move. Similar biases arise in the markets for labor and schooling, where some workers' and students' behaviors affect the quality of the environment experienced by other workers and students.⁸ In general, z_{2ht} should be treated as an endogenous variable that may change in response to a shift in z_{1ht} .

A similar form of bias arises if goods exist that are complements or substitutes with housing. One key interaction is between the markets for housing and labor. Consumers often decide where to live based upon job availability, and the types of jobs that are available in an area affect housing prices and the types of people who live there. Additionally, consumers who are considering buying a home in an area consider the local price level and the quality and variety of local goods. Hence, local job characteristics and local prices are location-specific attributes that should be included in Equation 1. In addition to being difficult to measure in an exhaustive way, these factors are probably correlated with ε_{ht} , and adequately controlling for these variables would require finding a credible instrument for each one.

The two-by-two matrix in Figure 1 illustrates how omitted attributes and products will bias OLS estimates of β_1 . The sign of the bias depends on the signs of its two components, $cov(z_{1ht}, z_{2ht})$ and β_2 . The first component, $cov(z_{1ht}, z_{2ht})$ indicates the extent to which an increase in z_{1ht} increases or decreases the supply of z_{2ht} . In the housing market, a positive improvement in one location-specific attribute like school quality might attract affluent and educated homebuyers and high-end businesses to the area. These changes would probably reduce crime, improve neighborhood cleanliness, and increase

⁷ If z_{2ht} is observable and data are available for some period preceding the shock to school quality, then one might consider instrumenting for z_{2ht} with the pre-treatment levels of that attribute. In general, however, this geographic variation in z_{2ht} will be correlated with many unobservable location-specific determinants of home value. Suppose, for example, that an additional omitted attribute z_{3ht} is the natural beauty of the area, which is time-invariant, positively correlated with z_{2ht} , and positively valued by consumers. Instrumenting for z_{2ht} with the pre-treatment value will produce an upward-biased estimate of β_2 that captures the effects of both pleasantness of neighbors (which changes in response to a shock to school quality) and natural beauty (which is time-invariant, does not respond to the shock, and should not appear in the set of controls). Supposing that $cov(z_{1ht}, z_{2ht})$, β_1 , and β_2 are all positive, this upward bias in the estimation of β_2 will generate a downward bias in the estimation of β_1 , so that the researcher attributes too much of z_{1ht} 's effects on housing prices to the increase in the pleasantness of neighbors.

⁸ Another problem that is particular to labor market studies is that individual wages are determined by workplace amenities and worker productivity, both of which are endogenous and difficult to measure.

Figure 1: Direction of Bias due to Omitted Attributes or Products


		Omitted attribute or product is a . . .	
		Good	Bad
Effect of attribute of interest on supply of omitted attribute or product	Positive	Positive	Negative
	Negative	Negative	Positive

the variety of local products. They would also increase congestion, prices, and tax rates. The second component, β_2 , indicates the extent to which z_{2ht} is a good or a bad in the typical consumer's utility function. If increasing school quality increases the variety of local products and product variety is a good, then failing to control for product variety creates a positive bias in the estimation of β_1 , as in the upper left square of the matrix. If a school quality improvement reduces crime and crime is a bad, then failing to control for crime also creates a positive bias, as in the lower right of the matrix. If increasing school quality increases congestion and the local price level and both are bads, then failing to control for those factors creates a downward bias, as in the upper right of the matrix.

The third source of bias, variation in aggregate quantity, is illustrated in Figure 2, which is adapted from Bartik (1987), who originally discussed this problem. The lower demand curve in Figure 2 plots consumer willingness to pay for living in that area before a school quality improvement. The higher demand curve shows willingness to pay for living in that area after the improvement. The benefit of the intervention is illustrated by the vertical difference in these two demand curves. The intervention's effect on prices, which is shown on the vertical axis, will in general be smaller than the benefit of the intervention. The price difference will only equal this benefit in the special case that supply is perfectly inelastic. In general, the price difference may equal MWTP in the short run, when quantity is fixed. In the long run, however, quantity will increase, and this price difference will shrink.

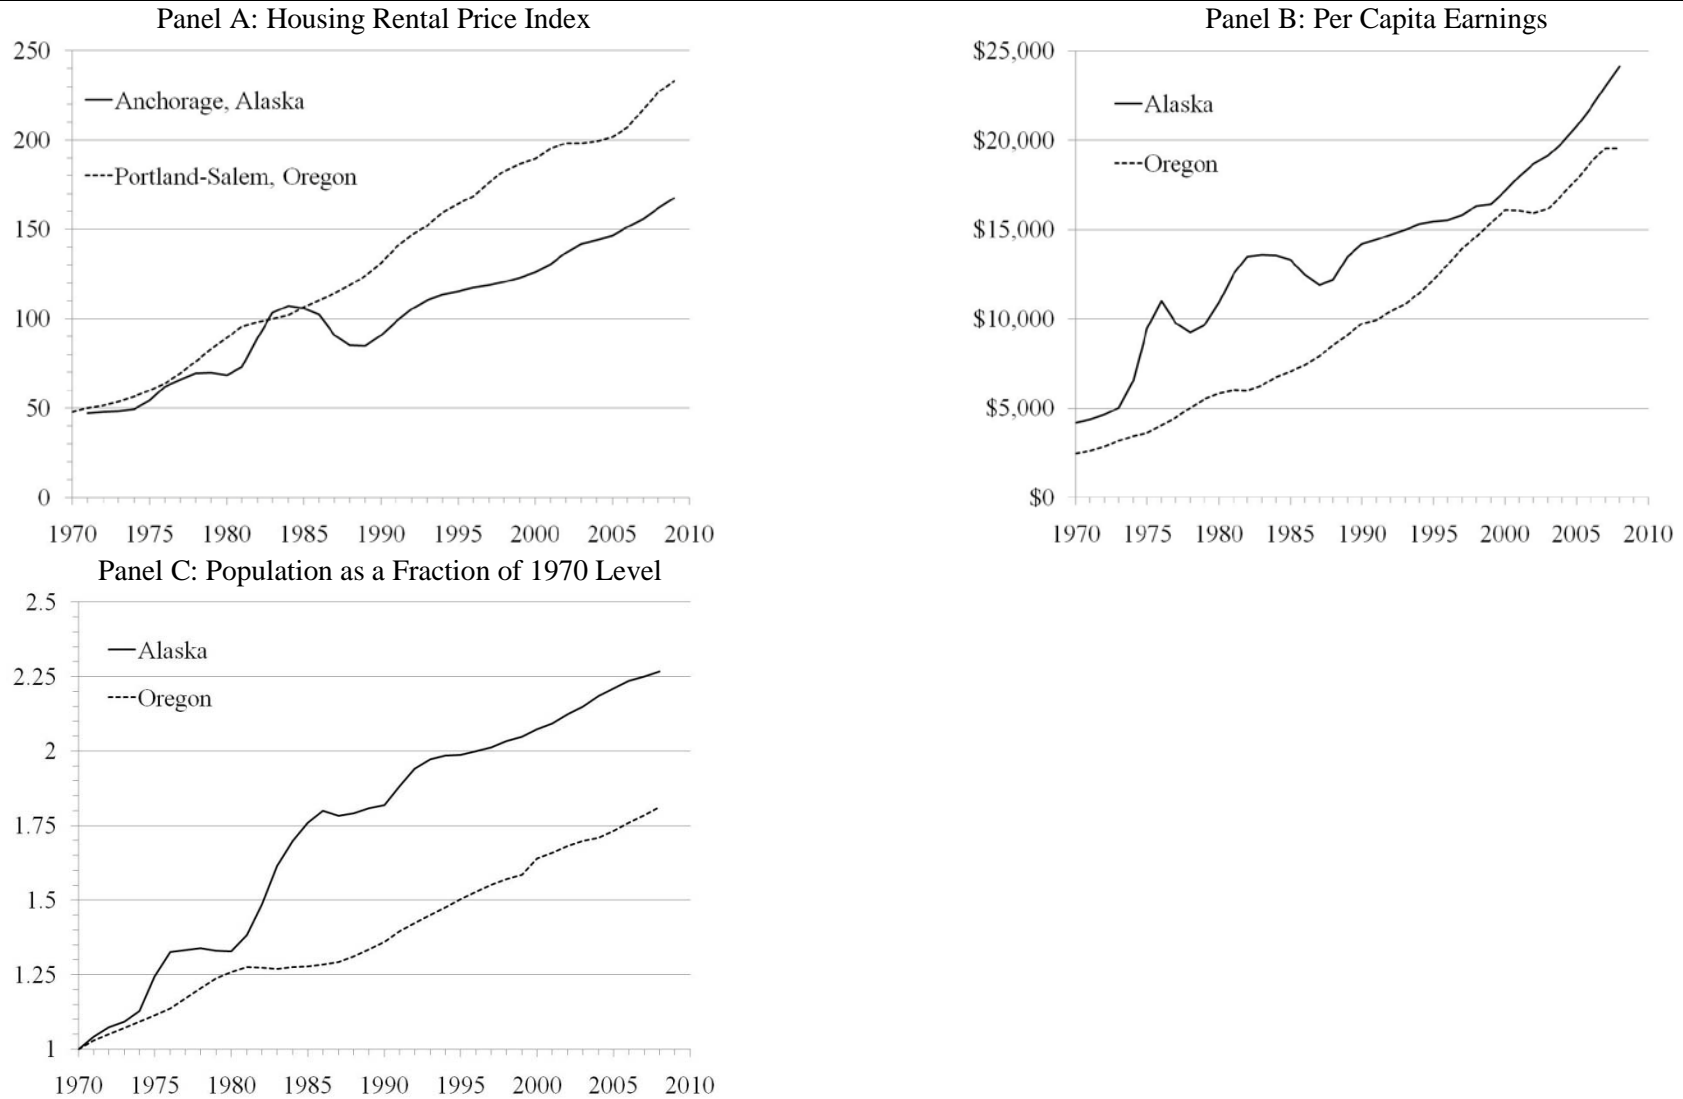
To illustrate the practical importance of these three sources of bias, panel A of Figure 3 shows indices of the rental price of housing in the Anchorage, Alaska and Portland-Salem, Oregon areas from 1970 to 2009, and panels B and C show per capita earnings and state population as fractions of the 1970 levels in Alaska and Oregon. In the 1970s, rent and wages increased due to a temporary rise in labor

Figure 2: Illustration of Bias due to Aggregate Quantity Changes



demand associated with oil pipeline construction (Carrington, 1996). They rose again in the 1980s due to growth in the oil industry and the 1982 initiation of annual cash payments to every resident (Alaska Permanent Fund Corporation, 2009; U.S. Energy Information Administration, 2010). Both increases underscore the extent to which wages and home prices are simultaneously determined. Measuring the total benefit of an oil boom or transfer program requires examining the effects on both rents and wages. The 1980s price increase was temporary and probably reflected the inability of housing and labor supply to respond quickly to the rises in demand. In the long run, population increased, as shown in panel C, and prices dropped back to trend. These long run quantity responses (and corresponding drops in prices) are assumed not to occur in traditional hedonic models. The oil jobs and cash transfers both made Alaska a more desirable place to live; however, these improvements did not generate long term price changes, and a standard hedonic comparison of housing prices between 1990 and 1980 would incorrectly indicate that these changes provided little or no value. Additionally, the influx of people to Alaska probably affected school quality, crime, the income, education, age, and sex composition of neighborhoods, the variety of

Figure 3: Housing Rental Prices, Per Capita Earnings, and Population in Alaska and Oregon, 1970-2009



Notes to Figure 3: Data for panel A taken from U.S. Bureau of Labor Statistics (2010); earnings data for panel B taken from U.S. Bureau of Economic Analysis (2010); population data for panels B and C taken from U.S. Bureau of the Census (2010).

local jobs, and the variety and prices of local products. Many of these changes are difficult to measure, and when they are measurable, it is unclear how one might control for them or combine these many responses into a single dollar value.

In the models of Rosen (1974) and Epple (1987), all of the benefits of a product improvement are capitalized into the price. In a more general framework in which supply is not perfectly inelastic, some of the benefits of the improvement will be capitalized into the price, and the remainder will be experienced as consumer surplus. If the affected product has complements or substitutes, the intervention could affect prices or production quantities for those other goods. Suppose, for example, that the supply of computers is perfectly competitive, but the supply of operating systems is controlled by a monopoly with inelastic supply. A technological innovation that increased computing speed could have no effect on the price of computers, but the producer of operating systems could raise its prices and keep output constant.

BLP (1995) develop a discrete choice model in which the dependent variable is the decision to purchase a specific good, and the price of the good is an endogenous regressor. The quantity effects described in Figure 2 and illustrated in Figure 3 are partially considered in their model, as an exogenous shock to product attributes can affect prices and market shares – *i.e.*, the fraction of sales accounted for by a single product – but not aggregate quantity produced. One limitation of BLP's approach is that, in addition to having an instrument for the product attribute being studied, it is necessary to have an instrument that shifts prices but is unrelated to unobserved factors influencing demand. The types of quasi-experiments that generate valid instruments are rare, and it would be very unusual to have a single dataset in which separate plausible instruments existed for both an important product attribute and the price of the good. Due to this restriction, researchers using the BLP approach have been forced to use instruments for the price that probably do not satisfy the necessary exogeneity conditions.

The current study addresses the problem of endogenous prices in two ways. For the first set of experimental estimators, the researcher offers products and subsidies to individual consumers. Offers made to these atomless consumers do not generate any general equilibrium responses in prices. For the second set of estimators developed in this study, it is necessary to have two samples: one in which there is exogenous variation in the product attribute being studied and another in which there is random variation in the price. The second sample is used to measure the slope of the demand curve for the product of interest. The estimated effects of the attribute on quantity consumed are used together with estimates of the slope of the demand curve to measure the vertical change in demand per additional unit of the attribute. For this vertical change to be identified, the slope of the demand curve must be similar across the two samples.

III. OVERVIEW OF ESTIMATORS

Table 1 describes the seven experimental and quasi-experimental estimation strategies that are developed in this study. These estimators illustrate what aspects of the demand for product attributes can be identified in different situations. Section V begins with an “idealized experiment” to highlight what assumptions are in principle necessary to identify MWTP. This new modeling concept serves as a theoretical benchmark to clarify the tradeoffs that researchers face between generality of the model and availability of experimental or quasi-experimental data. In the *idealized experiment*, the researcher has the ability to “treat” goods that a consumer purchases with an additional unit of the attribute of interest z_k . A small set of consumers is selected from the population, and each one is offered this treatment option at a different, randomly assigned price. Because this intervention only affects the small number of consumers participating in the study, it avoids biases from market-level changes to other attributes, other goods’ prices, or aggregate quantity. Through this *idealized experiment* it is possible to identify the distribution across consumers of the MWTP for the attribute z_k . In keeping with the policy focus of hedonic estimation, exclusive attention is paid to identification of the demand for product attributes, and extra assumptions are not imposed to identify utility or cost functions.

The remainder of Section V develops alternative, practical estimators that identify the distributions across consumers of policy-relevant measures of product attribute demand that differ slightly from MWTP. The second estimator involves an “existing tradeoff experiment.” To conduct this experiment, the researcher identifies an existing subset of the population in which each individual faces a tradeoff between an “untreated” set of goods and a “treated” version that has an additional unit of z_k . The third estimator uses an “offer-restricted experiment,” in which the researcher artificially restricts consumers’ options to the untreated and treated sets of goods and provides financial compensation to the participants to induce them to accept this restriction. In both of these experiments, the researcher identifies the value of the treatment by randomly assigning subsidies for selecting the treatment option across consumers in the study.

The estimators presented in Section VI compare total sales of untreated goods with their treated variants. For these research designs, some consumers face decisions between untreated goods and their best outside options, and others face decisions between treated goods and their outside options. Using these estimators, it is not possible to identify MWTP for any consumer, but it is possible to identify the effect of the attribute z_k on aggregate surplus, the vertical increase in demand illustrated in Figure 2. The fourth estimator in this study uses a “randomized offer experiment.” The researcher offers each participant in the experiment a subsidized good. The level of z_k in the offered good and the amount of the subsidy are randomized across consumers. The fifth estimator is based upon a “take-up study.” That study is similar to the *randomized offer experiment*, except that it involves two experiments: a “take-up experiment” and a “price experiment.” In the *take-up experiment*, the researcher makes subsidized

Table 1: Description of Estimators Developed in Study

	(1) Estimator	(2) Research design	(3) Key Assumptions	(4) Identifies . . .	(5) Applications described in text
1.	Idealized Experiment	Offer consumers the option to “treat” all of their goods in the market of interest with an additional unit of the attribute z_k . Randomize the price for the treated option across consumers	<i>Local non-satiation</i> , price-taking consumers.	Distribution across consumers of MWTP for an attribute z_k	Home improvements, product upgrades, and attributes that are artificially tied to houses or jobs, such as school district access or health care coverage.
2.	Existing Tradeoff Experiment	Identify consumers already making a decision between “treated” and “untreated” versions of the same good. Randomize subsidy for treated version across consumers.	Same as <i>idealized experiment</i> .	Distribution across subset of consumers facing tradeoff of MWTP for an attribute z_k	The discount rate, the value of freedom from jail, and the value of a statistical life.
3.	Restricted Offer Experiment	Pay consumers to restrict consumption in the market of interest to “treated” or “untreated” versions of the same good. Randomize price for treated version across consumers.	Same as <i>idealized experiment</i> .	Distribution across consumers of “offer-restricted MWTP” for an attribute z_k	The value of doctor visits and medical treatment, the value of internet bandwidth, the discount rate.
4.	Randomized Offer Experiment	Offer some consumers “untreated” goods other consumers “treated” goods, both at a subsidized rate, where the subsidy is randomly assigned across consumers.	Same as <i>idealized experiment</i> .	“Marginal Surplus” (i.e., the vertical difference between the treated and untreated demand curves) at all points along the demand curve.	The value of different characteristics of small business loans or solicitations for charitable donations.
5.	Take-Up Estimator	In a “take-up experiment,” offer some consumers “untreated” goods and other consumers “treated” goods, both at a subsidized rate. In a “price experiment,” randomly assign subsidies for untreated goods across consumers.	Same as <i>idealized experiment</i> plus <i>linear demand</i> and the slope of the demand curve is the same between the <i>take-up experiment</i> and the <i>price experiment</i> .	The Marginal Surplus for the marginal unit sold of the treated good.	Racial preferences by employers, banks offering loans, or home sellers. Parents' valuations of reductions in elementary school class sizes. The value of Section 8 vouchers (long-term rent subsidies).
6.	Placebo Consumers Estimator	There is an “aggregate treatment” that increases attribute z_k by one unit a certain subset of goods for all consumers. Measure consumption of goods in that subset before and after the aggregate treatment for both affected and placebo consumers. Additionally, conduct a price experiment that randomly assigns subsidies for untreated goods across consumers.	Same as <i>take-up estimator</i> plus there is a “placebo” group of consumers who are similar to other “affected” consumers but place no value on the attribute of interest. Affected and placebo consumers have equal responses to <i>aggregate treatment's</i> effects on prices and on other attributes.	The Marginal Surplus for the marginal unit sold of the treated good.	Parental valuations of changes in school quality, worker valuations of health care benefits packages, the value of avoiding the Vietnam draft, military recruits' valuations of GI Bill benefits.
7.	Combined Price and Quantity Estimator	There is an “aggregate treatment” that increases attribute z_k by one unit a certain subset of goods for all consumers. Measure consumption of goods in that subset before and after the aggregate treatment. Additionally, conduct a price experiment that randomly assigns subsidies for untreated goods across consumers.	Same as <i>take-up estimator</i> plus there are no endogenous amenities except z_k , and the only price that responds to the shock is that of the product being studied.	The Marginal Surplus for the marginal unit sold of the treated good.	Any previous hedonic study that only examines price effects. Recently applied as a specification check in a study by Rohlfs (2011b) measuring consumer valuations of automobile air bags.

product offers in which the level of z_k is randomized. In the price experiment, the researcher offers consumers products for which the subsidized price is randomized. For both the *randomized offer experiment* and the *take-up study*, the researcher compares the effect of z_k on the demand for the offered good to the effect of the dollar subsidy to measure the subsidy amount that would increase demand by as much as the attribute does. Using the logic of Figure 2, these approaches consider a case in which the price difference is constant (so that the supply curve is effectively flat) and use the horizontal quantity difference together with the slope of the demand curve to identify the vertical increase in demand resulting from the increase in the attribute. For the *take-up study*, it is necessary to assume that the slope of the demand curve for the good (and hence the effect of the subsidy) is similar between the *take-up experiment* and the *price experiment*.

The remaining estimators developed in Section VI examine quasi-experiments in which there is a shock that “treats” some goods for all consumers. The sixth approach, the “placebo consumers estimator” makes use of a *placebo group* of consumers who do not value the attribute of interest but are affected by the macroeconomic effects of the shock on endogenous omitted attributes and complements and substitutes for the good being studied. For example, if the shock is an improvement in school quality, non-parents might constitute a placebo group, as they are affected by the change in housing prices, but they do not value the treatment itself. A difference-in-differences design is used to disentangle the direct effects of the attribute of interest (which do not exist for placebo consumers) from the indirect effects on prices and other attributes (which do affect placebo consumers’ demand). The seventh approach, the “combined price and quantity study” considers a case in which there are no endogenous omitted attributes or complementary or substitute goods, so that the shock only affects the price and quantity demanded of the good being studied. The vertical increase in demand is estimated as the price change minus the slope of the demand curve times the quantity change. As with the *take-up study*, both the *placebo consumers study* and the *price and quantity comparison study* use data from the *price experiment* to convert the quantity effect into dollar units. Consequently, it is necessary in both studies to assume that the slope of the demand curve as estimated from the *price experiment* generalizes to these other contexts.

IV. CONCEPTUAL FRAMEWORK

A. The Consumption Decision

Each good in the economy is heterogeneous and is defined by its levels of different attributes.⁹ Let m be the number of attributes that define a good outside the market of interest, and let n denote the number of product attributes in the market of interest. The numbers m and n of attributes are large and describe every conceivable feature of every good. Let $\mathbf{X} \subset \mathbb{R}^m$ denote the set of all conceivable goods

⁹ In the literature, product attributes are also referred to as amenities or product characteristics.

outside the market of interest, and let $\mathbf{Z} \subset \mathbb{R}^n$ denote the set of all conceivable goods in the market of interest. Both \mathbf{X} and \mathbf{Z} are compact and continuous. A single good $\mathbf{x} \in \mathbf{X}$ or $\mathbf{z} \in \mathbf{Z}$ is a vector x_1, \dots, x_m or z_1, \dots, z_n of different attribute levels. If the market being studied is housing, then the attributes x_1, \dots, x_m are sufficiently detailed that they uniquely identify every good that is not a house, and the attributes z_1, \dots, z_n describe home- and location-specific features such as square footage and local air quality.

There is a mass $\mathbf{I} = [0,1]$ of atomless consumers, indexed by i . Let the non-negative scalars $q_{xi}(\mathbf{x})$ and $q_{zi}(\mathbf{z})$ denote the quantities of goods \mathbf{x} and \mathbf{z} that i consumes.¹⁰ Following Jones (1984), *consumption bundles* can be modeled as measures:

CONSUMPTION BUNDLE: A *consumption bundle* for consumer i consists of values $q_{xi}(\mathbf{x})$ and $q_{zi}(\mathbf{z})$ for every $\mathbf{x} \in \mathbf{X}$ and $\mathbf{z} \in \mathbf{Z}$. This consumption bundle is formally characterized as a pair (q_{xi}, q_{zi}) of finite, non-negative measures q_{xi} on \mathbf{X} and q_{zi} on \mathbf{Z} .¹¹

A typical measure on a set \mathbf{X} ascribes a number such as “distance from the origin” to each subset of \mathbf{X} , including single elements \mathbf{x} . The measures q_{xi} and q_{zi} are defined similarly; however, the numbers ascribed to each subset of \mathbf{X} and \mathbf{Z} are interpreted as “quantity consumed by i ” rather than distance from the origin.¹² Defining these measures provides a compact way to represent consumers’ complete consumption bundles. The term *consumption bundle* will describe the pair (q_{xi}, q_{zi}) , and each of the measures q_{xi} and q_{zi} will be referred to as a *bundle of goods*. Let the measure spaces \mathcal{M}_X and \mathcal{M}_Z contain all of the finite measures q_{xi} on \mathbf{X} and q_{zi} on \mathbf{Z} , and let $\mathcal{M}_{X[0,+]}$ and $\mathcal{M}_{Z[0,+]}$ be the subspaces of \mathcal{M}_X and \mathcal{M}_Z that contain only non-negative measures.

Let the scalars $p_x(\mathbf{x})$ and $p_z(\mathbf{z})$ denote the market prices per unit of \mathbf{x} and \mathbf{z} . Prices may be positive or negative. *Price schedules* are defined analogously to *consumption bundles*:

PRICE SCHEDULE: A *price schedule* consists of prices $p_x(\mathbf{x})$ and $p_z(\mathbf{z})$ for every $\mathbf{x} \in \mathbf{X}$ and $\mathbf{z} \in \mathbf{Z}$. Each price schedule is formally characterized as a pair (p_x, p_z) of signed and bounded measures p_x on \mathbf{X} and p_z on \mathbf{Z} .¹³

In principle, this definition of prices implies a thick market in which every good in \mathbf{X} and \mathbf{Z} can be purchased. As a practical matter, however, prices can vary discontinuously with the different attribute levels, and some prices may be very high, indicating a thin market in which those goods are effectively unavailable. The total cost of purchasing *consumption bundle* (q_{xi}, q_{zi}) is $\int_{\mathbf{X}} p_x(\mathbf{x}) * q_{xi}(\mathbf{x}) d\mathbf{x} +$

¹⁰ A static setting is considered here; however, it is straightforward to incorporate dynamics into the model by making “time in which the good is consumed” be one of the product attributes for both \mathbf{X} and \mathbf{Z} .

¹¹ All measures and subsets in this paper are defined in the Borel sense.

¹² For subsets of goods $\mathbf{X}_T \subseteq \mathbf{X}$ and $\mathbf{Z}_T \subseteq \mathbf{Z}$, the scalars $q_{xi}(\mathbf{X}_T) = \int_{\mathbf{X}_T} q_{xi}(\mathbf{x}) d\mathbf{x}$ and $q_{zi}(\mathbf{Z}_T) = \int_{\mathbf{Z}_T} q_{zi}(\mathbf{z}) d\mathbf{z}$ denote the total quantities of goods within those sets that i consumes.

¹³ These measures are defined so that the scalars $p_x(\mathbf{X}_T) = \int_{\mathbf{X}_T} p_x(\mathbf{x}) d\mathbf{x} / \int_{\mathbf{X}_T} 1 d\mathbf{x}$ and $p_z(\mathbf{Z}_T) = \int_{\mathbf{Z}_T} p_z(\mathbf{z}) d\mathbf{z} / \int_{\mathbf{Z}_T} 1 d\mathbf{z}$ denote the average prices of goods within $\mathbf{X}_T \subseteq \mathbf{X}$ or $\mathbf{Z}_T \subseteq \mathbf{Z}$

$\int_{\mathbf{Z}} p_z(\mathbf{z}) * q_{zi}(\mathbf{z}) d\mathbf{z}$. For simplicity (and with apologies for the abuse of notation), the costs of q_{xi} and q_{zi} are denoted $p_x \cdot q_{xi}$ and $p_z \cdot q_{zi}$.

Each consumer is endowed with a consumption bundle $(\omega_{xi}, \omega_{zi})$, with $\omega_{xi} \in \mathcal{M}_{X[0,+]}$ and $\omega_{zi} \in \mathcal{M}_{Z[0,+]}$. Let the scalar w_i denote i 's wealth $p_x \cdot \omega_{xi} + p_z \cdot \omega_{zi}$. The distribution of wealth for the economy can be expressed as follows:

WEALTH MEASURE. The *wealth measure for the economy* is a non-negative, finite measure $w_{\mathbf{I}}$ on \mathbf{I} that describes the wealth $w_{\mathbf{I}}(i) = w_i$ for each consumer i .

To simplify the notation, let prices and the distribution of wealth together be summarized by a single measure, denoted the *state of the economy*:

STATE OF THE ECONOMY. The state $s = (p_x, p_z, w_{\mathbf{I}})$ of the economy consists of a *price schedule* (p_x, p_z) and a *wealth measure* $w_{\mathbf{I}}$.

Let $\mathcal{M}_{\mathbf{I}[0,+]}$ denote the space of non-negative, finite measures on \mathbf{I} , and let \mathcal{M}_s denote the metric space $\mathcal{M}_X \times \mathcal{M}_Z \times \mathcal{M}_{\mathbf{I}[0,+]}$ of possible states of the economy.

Each consumer i has a utility function $u_i: \mathcal{M}_{X[0,+]} \times \mathcal{M}_{Z[0,+]} \rightarrow \mathbb{R}$. This function $u_i(q_{xi}, q_{zi})$ depends upon i 's consumption levels $q_{xi}(\mathbf{x})$ and $q_{zi}(\mathbf{z})$ for each good $\mathbf{x} \in \mathbf{X}$ and $\mathbf{z} \in \mathbf{Z}$. Consumers take prices as given.¹⁴ Each consumer's utility function is *locally non-satiable* – i.e., utility strictly increases with wealth. Utility may be non-separable across goods and non-monotonic and discontinuous in any of the goods.¹⁵ Utility functions may vary in arbitrary ways across consumers; however, the quantities $q_{xi}(\mathbf{x})$ and $q_{zi}(\mathbf{z})$ are restricted to be measurable across agents for all \mathbf{x} and \mathbf{z} , with the economy-wide levels of consumption of \mathbf{x} and \mathbf{z} represented by $\int_{\mathbf{I}} q_{xi}(\mathbf{x}) di$ and $\int_{\mathbf{I}} q_{zi}(\mathbf{z}) di$ (cf., Aumann, 1964).

Consumer i selects a *consumption bundle* (q_{xi}, q_{zi}) to maximize u_i subject to the constraint that $p_x \cdot q_{xi} + p_z \cdot q_{zi} \leq w_i$. The remainder of this paper assumes the existence of equilibrium prices and focuses on the identification of consumer demand for product attributes. Let the equilibrium price schedule be denoted (p_x^*, p_z^*) . Let i 's wealth $p_x^* \cdot \omega_{xi} + p_z^* \cdot \omega_{zi}$ at these prices be denoted w_i^* , and let s_i^* denote the state of the economy at these equilibrium prices. Let (q_{xi}^*, q_{zi}^*) denote the consumption bundle that solves i 's constrained maximization problem given market prices (p_x^*, p_z^*) , and let the scalar u_i^* denote $u_i(q_{xi}^*, q_{zi}^*)$, the level of utility provided by i 's optimal *consumption bundle*. Supply may be competitive or non-competitive. In practice, additional assumptions about the market must be imposed if one wishes to ensure that an equilibrium exists. Production and equilibrium are discussed in the appendix.

¹⁴ The restriction that consumers are price-taking is non-trivial. Ostroy (1984) shows that, when there is a continuum of atomless goods, atomless consumers can have market power. Such market power is ruled out here.

¹⁵ Additionally, the convexity or concavity of u_i in each good \mathbf{z} may vary with $q_{zi}(\mathbf{z})$. A sufficient but not necessary condition for this utility function to exist is that the preference relation is continuous in all good quantities and attributes, as is formally stated in Jones (1984, pg. 511).

Previous hedonic and discrete choice models by Rosen (1974), Epple (1987), and BLP (1995) are special cases of the framework described here. All of those models restrict the space \mathbf{X} to be a single point, so that the commodity space consists of \mathbf{Z} and a homogeneous consumption good. In the housing market, a minor extension to the conventional models would allow \mathbf{X} to consist of two homogeneous goods, those purchased in Atlanta and those purchased in Boston. A home buyer would certainly take the price level into account when choosing where to live, and a change in weather or school quality in one of those cities could affect local prices for the homogeneous good as well as local housing prices. The full market response to the weather or school quality shock includes both the effect on housing prices and the effect on local prices for everything else. These local price effects, which are assumed not to exist in previous hedonic and discrete choice models, are easy to accommodate in the current setup.

Another important type of complementarity occurs between jobs and housing in the same location. A firm hiring in Syracuse, New York (which receives 115 inches of snow per year) must pay a higher wage to obtain the same level of talent than does a similar firm hiring in San Francisco (which has year-round pleasant weather and is located on the coast). In the current model, one of the goods \mathbf{x} could be hours of work in a specific job in Syracuse, and another good \mathbf{x}' could be hours of work in a specific job in San Francisco. Each consumer would be endowed with quantities of \mathbf{x} and \mathbf{x}' , which could be sold to the employer or consumed as leisure. The shape of the utility function would be such that selling hours to the firm in San Francisco (and failing to consume them as leisure) would greatly increase the utility benefit of housing in San Francisco. The goods \mathbf{x} and \mathbf{x}' would also be strong substitutes in the utility function, so that selling hours of work in San Francisco would sharply increase the utility cost of selling hours of work in Syracuse, and working jobs in both cities would be rare.

In addition to ruling out substitutes and complements, previous hedonic models require that each consumer purchases exactly one unit of one good \mathbf{z} and zero units of every other good in \mathbf{Z} . In the market for housing, this restriction rules out any effects of prices or location-specific attributes on the decision to buy a home or the total number of home buyers.¹⁶ In addition to allowing for endogenous homeownership, relaxing this restriction helps to describe markets in which consumers often buy more than one good, such as automobiles, computers, and even markets in which quantity consumed varies continuously, as in foods, music, and vacations.

¹⁶ Additionally, Rosen and BLP assume that utility is a monotonic function of each attribute for all consumers and that production costs increase in each attribute for all firms. Rosen also assumes concave utility, convex costs, observability of every attribute, “thick” markets, and perfectly competitive firms. Epple uses the same set of assumptions but allows for the existence of unobservable attributes, provided that they are uncorrelated with geographic or temporal variation in observed attributes. BLP also assume that utility and log costs are linear, preference heterogeneity follows an exponential distribution, the one unobserved attribute is uncorrelated with observed ones and is valued the same by all buyers, and firms are oligopolists.

B. The Bid Function, Reservation Prices, and Consumer Surplus

Following Rosen (1974), MWTP is expressed in terms of a bid function, reservation prices, and consumer surplus:

BID FUNCTION: *Consumer i 's bid function* $\theta_i: \mathcal{M}_{Z[0,+]} \times \mathcal{M}_X \times \mathbb{R}^2 \rightarrow \mathbb{R}$ gives the price $\theta_i(q_Z, p_X, w, u)$ for a bundle of goods q_Z at which consuming that bundle and spending all remaining wealth optimally on goods in \mathbf{X} produces utility level u . Taking q_Z, p_X, w and u as given and letting λ_i denote consumer i 's Lagrange multiplier, the *bid function* returns the value of θ that satisfies $\max_{q_{xi}} \{u_i(q_{xi}, q_Z) + \lambda_i * [w - p_X \cdot q_{xi} - \theta]\} = u$.

RESERVATION PRICE: *Consumer i 's reservation price for the bundle q_Z* equals $\theta_i(q_Z, p_X^*, w_i^*, u_i^*)$. This scalar measures the price for q_Z at which i would be indifferent between purchasing q_Z or purchasing q_{zi}^* at the price $p_Z^* \cdot q_{zi}^*$.

CONSUMER SURPLUS: The *consumer surplus to i from consuming q_Z* equals $\theta_i(q_Z, p_X^*, w_i^*, u_i^*) - p_Z^* \cdot q_Z$. This scalar measures the excess value obtained from consuming q_Z at the price $p_Z^* \cdot q_Z$ rather than consuming q_{zi}^* at the price $p_Z^* \cdot q_{zi}^*$.

Local non-satiation guarantees that, for each q_Z, p_X and w , there is no more than one value of θ at which i 's utility equals u . If two prices θ and θ' produced the same utility level, with $\theta < \theta'$, then a cash subsidy of $\theta' - \theta$ would have no effect on utility, thus violating *local non-satiation*. Some values of u may not be achievable for a given bundle of goods q_Z due to a jump discontinuity in the utility function. In such cases, no θ solves the condition above, and θ is defined to equal the price of q_Z at the discontinuity point. Some bundles q_Z may be so undesirable that, if i consumes q_Z , no cash subsidy could increase utility to as high a level as u . In these cases, θ is defined to equal ∞ . Other bundles may be so desirable that, if consumed, no tax could reduce utility to a level as low as u . In these cases, θ is defined to equal $-\infty$.

In a standard model of a homogeneous good, consumer i 's *reservation price* for the good is the price at which i would be indifferent to purchasing the good or not. If this definition were used in the current context, i 's *reservation price* for the bundle q_Z would be evaluated at prices p_X^* , wealth w_i^* , and the level of utility that i would obtain if i consumed zero goods in \mathbf{Z} and spent all wealth optimally on goods in \mathbf{X} . In the hedonic context, defining *reservation prices* in this way is problematic. For many heterogeneous goods such as housing, food, or medical care, it is not natural to think about what utility would be if no goods in \mathbf{Z} were consumed. For this reason and following Rosen (1974), *reservation prices* and *consumer surplus* are defined using u_i^* as a benchmark utility level. Consumer i 's *reservation price* $\theta_i(q_{zi}^*, p_X^*, w_i^*, u_i^*)$ for the bundle q_{zi}^* is necessarily $p_Z^* \cdot q_{zi}^*$, and the *consumer surplus* to i from consuming q_{zi}^* is necessarily zero.

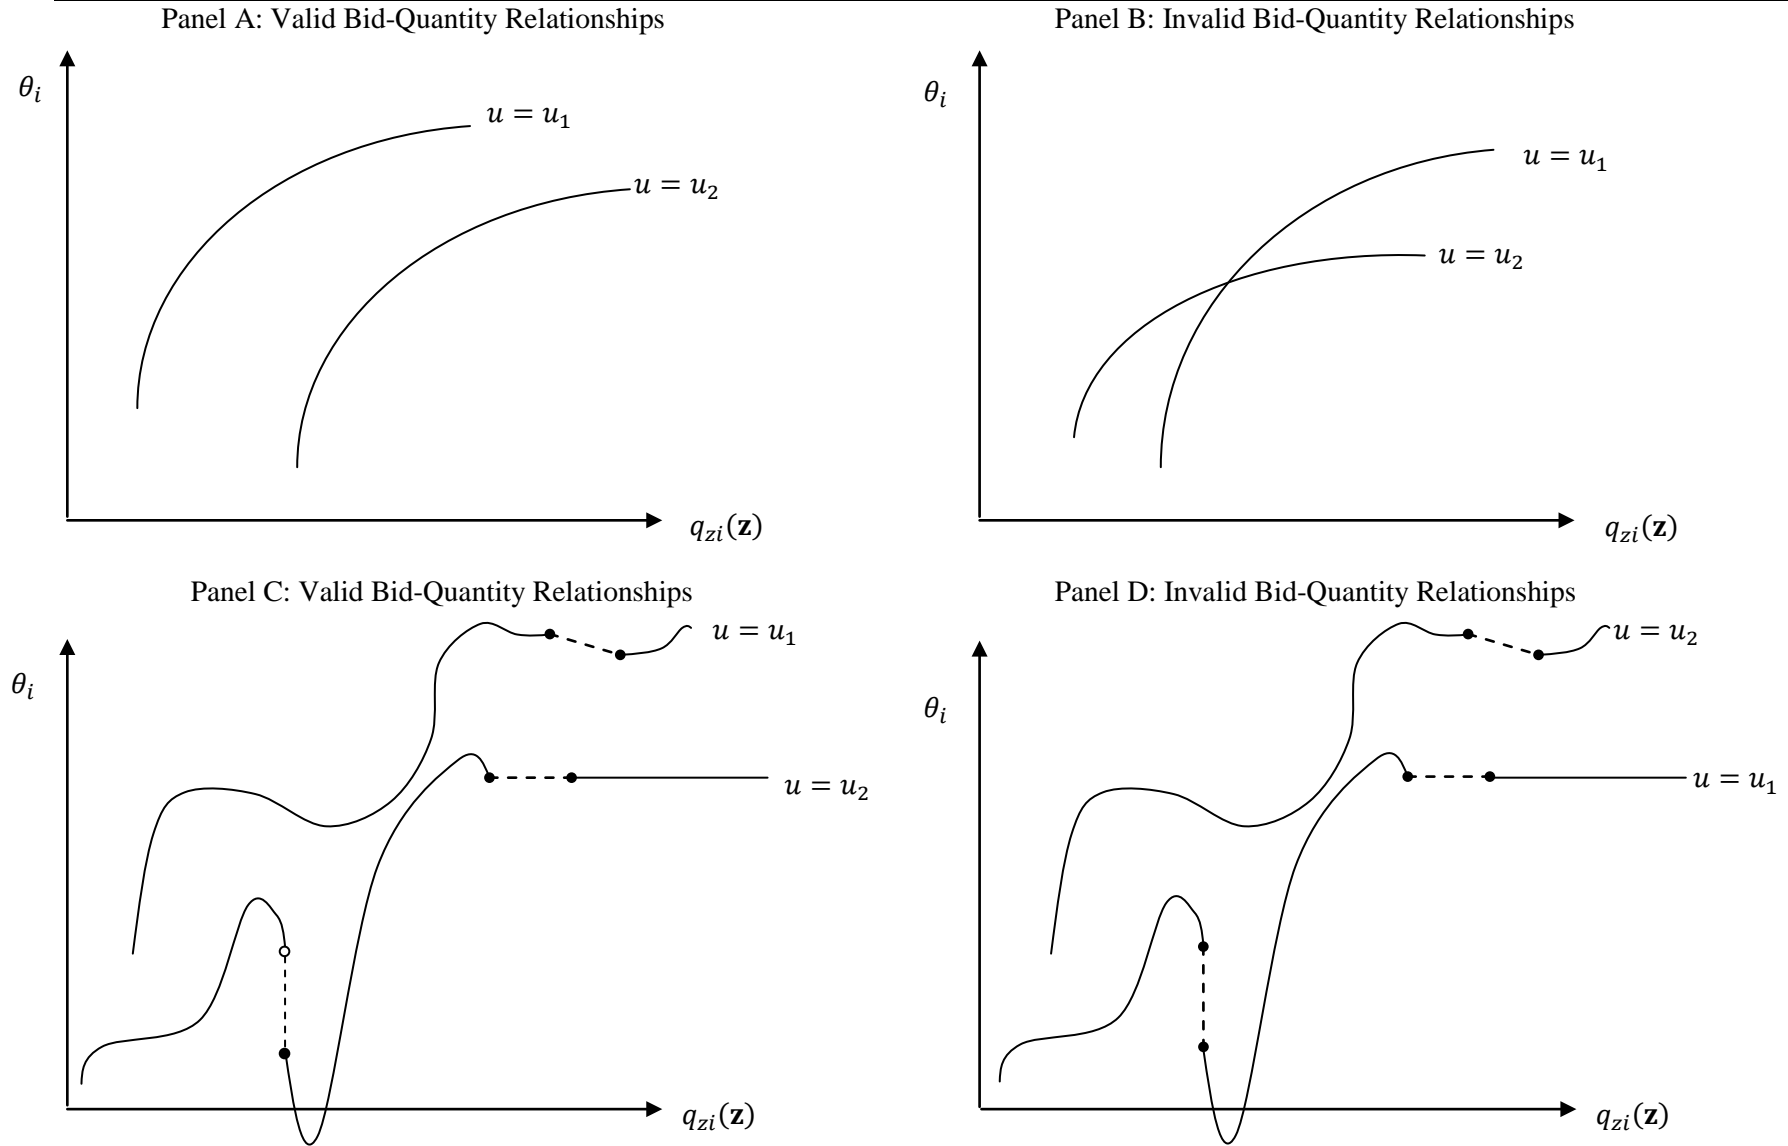
For a specific good \mathbf{z} , the relationship between θ_i and $q_{zi}(\mathbf{z})$ maps out an indifference curve showing the tradeoffs between dollars and units of \mathbf{z} consumed. The utility function may have an unusual

shape, but *local non-satiation* imposes some restrictions on how these indifference curves can look. Figure 4 shows examples of plots of θ_i against $q_{zi}(\mathbf{z})$ for a single consumer. Panels A and C show bid-quantity relationships that are allowed by the model, and panels B and D show bid-quantity relationships that violate the assumptions of the model. Panel A shows a typical pair of convex indifference curves. Utility is a decreasing function of the amount paid for q_z , and $q_{zi}(\mathbf{z})$ is a good, so that utility is lowest in the upper left of the graph and highest in the lower right. In panel B, the indifference curves cross. At this crossing point, a single price θ_i and quantity $q_{zi}(\mathbf{z})$ produce two different utility values u_1 and u_2 , which violates the assumption that utility is a scalar-valued function. The indifference curves in panel C are non-monotonic, take on negative values, have a jump discontinuity, and equal ∞ or $-\infty$ for some values of $q_{zi}(\mathbf{z})$. The dashed horizontal and diagonal portions of the curves show $q_{zi}(\mathbf{z})$ values at which θ_i is positive or negative infinity. Some of the undefined values appear on the lower curve. At these levels of $q_{zi}(\mathbf{z})$, it is possible to achieve a utility level of u_1 , but no cash payment could increase i 's utility to as high as u_2 . In order to satisfy *local non-satiation* and to avoid utility functions crossing, the *bid function* must also be undefined over the same range of $q_{zi}(\mathbf{z})$ for all utility levels greater than u_2 . For the $q_{zi}(\mathbf{z})$ values at which the higher curve has dashes, there is no tax that could reduce i 's utility to as low as u_1 . For *local non-satiation* to be satisfied, the *bid function* must also be undefined over the same range of $q_{zi}(\mathbf{z})$ for all utility levels lower than u_1 . Panel D shows the same curves as in panel C but with two changes, both of which violate *local non-satiation*. First, at the jump discontinuity, there is a single $q_{zi}(\mathbf{z})$ at which the bid function takes on two values. Second, utility is higher for the higher curve than for the lower one, so that raising the cost of q_z increases utility.

In Rosen's (1974) original study, each consumer selects only one unit of one good \mathbf{z} , and the *bid function* depends upon which \mathbf{z} is chosen. Rosen imposes sufficient conditions to guarantee that this function is smooth and convex in each attribute z_k . Similar bid-attribute curves could be drawn in the current setting if q_z is restricted to equal one for a single good \mathbf{z} and zero for all other goods. In that case, the four panels of Figure 4 could be redrawn with z_k rather than $q_{zi}(\mathbf{z})$ along the horizontal axis. In Rosen's framework, only panel A is a valid bid-attribute relationship, and the slope of the bid-attribute relationship must have the same sign for every consumer. Using the current model, panels A and C are both valid, and the relationship may vary in arbitrary ways across consumers.¹⁷ Allowing for such complex functional forms helps to accommodate product attributes such as sweetness, brightness, and number of garages, that are only desirable in small doses and others, such as proximity to specific friends and family, that only matter to certain consumers.

¹⁷ For BLP (1995), who assume linear utility, the only valid bid-attribute relationships would be straight lines.

Figure 4: Examples of Valid and Invalid Bid-Quantity Relationships



Notes to Figure 4: Each graph shows hypothetical plots of $\theta_i(q_z, p_x, w, u)$ against $q_{zi}(\mathbf{z})$ for different utility levels. In all cases, $u_1 < u_2$.

C. Treatment Technology and Marginal Willingness to Pay

The MWTP by consumer i for attribute z_k is defined in terms of a *treatment technology*. The treatment technology is defined as follows:

TREATMENT TECHNOLOGY. A *treatment technology* for attribute z_k and set \mathbf{Z}_T is a function $T_{\mathbf{Z}_T k}: \mathcal{M}_{\mathbf{Z}[0,+]} \rightarrow \mathcal{M}_{\mathbf{Z}[0,+]}$ that satisfies $q'_z = T_{\mathbf{Z}_T k}(q_z) \Leftrightarrow q'_z(\mathbf{z}) = q_z(z_1, \dots, z_k - 1, \dots, z_n)$ for all $\mathbf{z} \in \mathbf{Z}_T$ and $q'_z(\mathbf{z}) = q_z(\mathbf{z})$ for all $\mathbf{z} \notin \mathbf{Z}_T$.

Hence, $T_{\mathbf{Z}_T k}$ converts any bundle q_z of goods in \mathbf{Z} into a new bundle $T_{\mathbf{Z}_T k}(q_z)$ of goods in \mathbf{Z} in which every unit of every good in \mathbf{Z}_T is replaced with the equivalent good with an additional unit of z_k . This *treatment technology* effectively increases z_k by one unit for every \mathbf{z} in \mathbf{Z}_T that i chooses to consume. Applying the *treatment technology* to a positive mass \mathbf{Z}_T of goods for a positive mass of consumers could in principle generate changes in the attributes of other goods or in the prices of goods in \mathbf{Z} or \mathbf{X} . For the purposes of defining MWTP, we consider an application of the *treatment technology* holding these factors constant.

Knowing that the treatment technology will be applied to one's bundle of goods alters a consumer's optimal choices. Under the regime in which a treatment technology for attribute z_k will be applied to i 's goods in the set \mathbf{Z}_T , holding prices constant, i solves the following problem:

$$(1) \max_{q_{xi}, q_{zi}} u_i(q_{xi}, T_{\mathbf{Z}_T k}(q_{zi})) \text{ subject to } p_x \cdot q_{xi} + p_z \cdot q_{zi} \leq w_i.$$

Let (q_{xi}^*, q_{zi}^*) continue to denote i 's optimal consumption bundle at equilibrium prices in the absence of any intervention. Let (q_{xi}^k, q_{zi}^k) denote the consumption bundle that i would choose to purchase at equilibrium prices, given the knowledge that the bundle of goods q_{zi}^k will be converted into $T_{\mathbf{Z}_T k}(q_{zi}^k)$, where the treatment is applied to the entire set \mathbf{Z} . If, for example, z_k is local school quality, then i might select a home in a relatively low quality school district with the understanding that the district will be improved by the *treatment technology*. The bundle q_{zi}^k includes the home in the low quality district, and $T_{\mathbf{Z}_T k}(q_{zi}^k)$ represents the same bundle after the school quality improvement.

Given this technology, we define the benefit to i of a one-unit increase in attribute z_k :

MARGINAL WILLINGNESS TO PAY. Consumer i 's *Marginal Willingness to Pay (MWTP)* for z_k is scalar-valued, is denoted $MWTP_{ki}$, and equals $\theta_i(T_{\mathbf{Z}_T k}(q_{zi}^k), p_x^*, w_i^*, u_i^*) - p_z^* \cdot q_{zi}^k$.

$MWTP_{ki}$ is a dollar-denominated measure of the *consumer surplus* that i experiences due to the *treatment technology*, i.e., from consuming $T_{\mathbf{Z}_T k}(q_{zi}^k)$ at the price $p_z^* \cdot q_{zi}^k$. Because consumer surplus is defined relative to the benchmark utility level u_i^* , this formula gives the change in surplus that i would experience from switching from the optimal untreated bundle (which provides zero surplus) to the optimal treated bundle.

In order for $MWTP_{ki}$ to accurately measure this benefit, it is essential that the final term, $p_z^* \cdot q_{zi}^k$, be subtracted off, so that the expression returns a surplus and not a reservation price. If i lives in a school district with high housing prices and z_k is school district quality, being offered the treatment technology could induce i to move to a more affordable area. After applying the treatment technology, the new area could be as desirable as or even less desirable than the old one. However, the treatment technology had a positive benefit by helping i to save money. This savings is taken into account by subtracting the price difference, $p_z^* \cdot (q_{zi}^k - q_{zi}^*)$, which is negative in this case, from the difference in reservation prices.

In Rosen's (1974) original study, the MWTP is the derivative of θ_i with respect to z_k for the one good in \mathbf{Z} that i consumes. Because the market is "thick" with every possible good in \mathbf{Z} , each consumer selects an interior solution for each attribute, and the MWTP for z_k is exactly equal to the marginal cost to producers of adding an additional unit of z_k to a good. In that framework, all owners of Ford Mustangs place the same value on additional units of safety, horsepower, and fuel efficiency. In the current setting, prices are not necessarily continuous functions of product attributes, markets are thin, and two buyers of Ford Mustangs may have very different utility functions. Each might select a Mustang with factory specifications because it is one of the few products available at a particularly reasonable price.

Because we use a flexible functional form for u_i and because markets may be thin, i 's consumption bundles will not necessarily represent interior solutions. Consequently, it is more natural to consider discrete changes in attributes (as achieved through the treatment technology) than to differentiate the bid function. The discrete change of the treatment technology is also better than marginal changes are at reflecting the type of variation that is generated through experiments and quasi-experiments. Because it is possible to consume more than one \mathbf{z} in the current model, it is necessary in the definition of MWTP to specify which goods receive the additional unit of z_k . The definition of MWTP used here supposes that the *treatment technology* is applied to all goods in \mathbf{Z} that i consumes.

D. Heterogeneity in MWTP

The variation across consumers in MWTP can be modeled in terms of a distribution:

MWTP DENSITY. Let the *MWTP density for attribute z_k* be denoted $f_k^{MWTP}: \mathbb{R} \rightarrow \mathbb{R}_{[0,+]}$ and represent the population density function (pdf) of $MWTP_{ki}$ across consumers. Let $F_k^{MWTP}: \mathbb{R} \rightarrow [0,1]$ denote the Cumulative Distribution Function (CDF) of f_k^{MWTP} .

The *MWTP density for attribute z_k* describes the variation across consumers in their excess valuations of $T_k(q_{zi}^k, \mathbf{Z})$, their optimal bundles of goods in \mathbf{Z} in the case in which the *treatment technology* for z_k is applied to every good in \mathbf{Z} .

A primary goal of this study is to establish conditions for identifying the MWTP density. In many cases, ethical or cost considerations will prevent researchers from conducting experiments to accomplish this goal. To address those cases, experiments and quasi-experiments are considered that estimate alternative measures of attribute demand that are similar to MWTP.

E. Aggregate Demand and Marginal Surplus

Many recent hedonic studies examine the effects of shocks that increase the levels of some product attributes for entire communities. Chay and Greenstone (2005), for example, measure the effects of county-wide air pollution reductions on housing prices. When evaluating quasi-experiments of this form, it is useful to have a definition of the aggregate benefits of a large scale shock. Let $\mathbf{a} \subseteq \mathbf{I}$ be a subset of consumers, and let $\mathbf{Z}_T \subset \mathbf{Z}$ be a subset of goods in \mathbf{Z} . Let τ be a tax that any consumer in \mathbf{a} must pay per unit of goods consumed in \mathbf{Z}_T . Our definition of aggregate benefits depends upon the *aggregate demand function*:

AGGREGATE DEMAND FUNCTION. The *aggregate demand function for \mathbf{Z}_T among consumers in \mathbf{a}* is denoted $Q_{\mathbf{Z}_T\mathbf{a}}: \mathbb{R} \times \mathcal{M}_s \rightarrow \mathbb{R}_{[0,+]}$. For every i in \mathbf{a} , let q'_{zi} be the bundle of goods in \mathbf{Z} that is derived from solving the following problem: $\max_{q'_{xi}, q'_{zi}} \left\{ u_i(q'_{xi}, q'_{zi}) + \lambda_i * \left[w_i - p_x \cdot q'_{xi} - p_z \cdot q'_{zi} - \tau * \int_{\mathbf{Z}_T} 1 dz \right] \right\}$. The quantity $Q_{\mathbf{Z}_T\mathbf{a}}(\tau, s)$ is defined as $\int_{\mathbf{a}} \int_{\mathbf{Z}_T} q'_{zi}(\mathbf{z}) dz di$.

This function measures the total quantity of goods in \mathbf{Z}_T demanded by consumers in \mathbf{a} . It depends on the tax on goods in \mathbf{Z}_T and the state of the economy s . Because this function describes the demand multiple goods together, there is not a single ‘‘own price’’ of goods in \mathbf{Z}_T . The tax τ is introduced into the *aggregate demand function* to approximate an own price effect, because a change in τ has equal effects on the purchase prices for all goods in \mathbf{Z}_T . For simplicity, suppose that, in the absence of any intervention, τ is zero. In addition to measuring totals across consumers, the *aggregate demand function* can apply to an individual consumer i , where $Q_{\mathbf{Z}_T i}(\tau, s)$, denotes the total quantity of goods in \mathbf{Z}_T that would be consumed by i given a tax τ and state of the economy s . This function can also apply to a single good \mathbf{z} , where $Q_{\mathbf{z}\mathbf{a}}(\tau, s)$ refers to the total quantity of the good \mathbf{z} demanded by consumers in \mathbf{a} .

Product quantities vary continuously, and aggregate demand is a twice continuously differentiable and bijective function of τ for any set of goods $\mathbf{Z}_T \subseteq \mathbf{Z}$ and any consumers $\mathbf{a} \subseteq \mathbf{I}$. Given this bijective property, we can define the inverse of aggregate demand:

INVERSE AGGREGATE DEMAND. The *inverse aggregate demand function for \mathbf{Z}_T among consumers in \mathbf{a}* is denoted $\tau_{\mathbf{Z}_T\mathbf{a}}: \mathbb{R}_{[0,+]} \times \mathcal{M}_s \rightarrow \mathbb{R}$. This function $\tau_{\mathbf{Z}_T\mathbf{a}}(Q, s)$ is defined as the value of τ that satisfies $Q_{\mathbf{Z}_T\mathbf{a}}(\tau, s) = Q$.

At a given aggregate quantity level Q and state of the economy s the inverse aggregate demand function gives the positive or negative tax that would be required to make aggregate demand for \mathbf{Z}_T among consumers in \mathbf{a} exactly equal to Q . This tax can be interpreted as the surplus obtained from the Q^{th} unit of \mathbf{Z}_T demanded by consumers in \mathbf{a} .

Next, consider an exogenous one-unit increase in z_k for all goods in \mathbf{Z}_T , holding all prices constant. An example of such a shock could be a regulation that improves air quality in a neighborhood or one that requires certain automobiles to have air bags.

AGGREGATE TREATED DEMAND. The *aggregate z_k -treated demand function for \mathbf{Z}_T among consumers in \mathbf{a}* is denoted $Q_{\mathbf{Z}_T\mathbf{a}}^k : \mathbb{R} \times \mathcal{M}_s \rightarrow \mathbb{R}_{[0,+]}$. For every i in \mathbf{a} , let q'_{zi} denote the bundle of goods in \mathbf{Z} that is derived from solving the following problem: $\max_{q'_{xi}, q'_{zi}} \left\{ u_i(q'_{xi}, T_{\mathbf{Z}_T k}(q'_{zi})) + \lambda_i * [w_i - p_x \cdot q'_{xi} - p_z \cdot q'_{zi} - \tau * \int_{\mathbf{Z}_T} 1 d\mathbf{z}] \right\}$. The quantity $Q_{\mathbf{Z}_T\mathbf{a}}^k(\tau, s)$ is defined as $\int_{\mathbf{a}} \int_{\mathbf{Z}_T} q'_{zi}(\mathbf{z}) dz di$.

INVERSE AGGREGATE TREATED DEMAND. The *inverse aggregate z_k -treated demand function for \mathbf{Z}_T among consumers in \mathbf{a}* is denoted $\tau_{\mathbf{Z}_T\mathbf{a}}^k : \mathbb{R}_{[0,+]} \times \mathcal{M}_s \rightarrow \mathbb{R}$. This function $\tau_{\mathbf{Z}_T\mathbf{a}}^k(Q, s)$ is defined as the value of τ that satisfies $Q_{\mathbf{Z}_T\mathbf{a}}^k(\tau, s) = Q$.

The *aggregate z_k -treated demand function* measures the total quantity $\int_{\mathbf{a}} \int_{\mathbf{Z}_T} q_{zi}(\mathbf{z}) dz di$ of goods in \mathbf{Z}_T that would be demanded by consumers in \mathbf{a} if consumers knew beforehand that the treatment technology for z_k would be applied to all goods purchased in \mathbf{Z}_T .

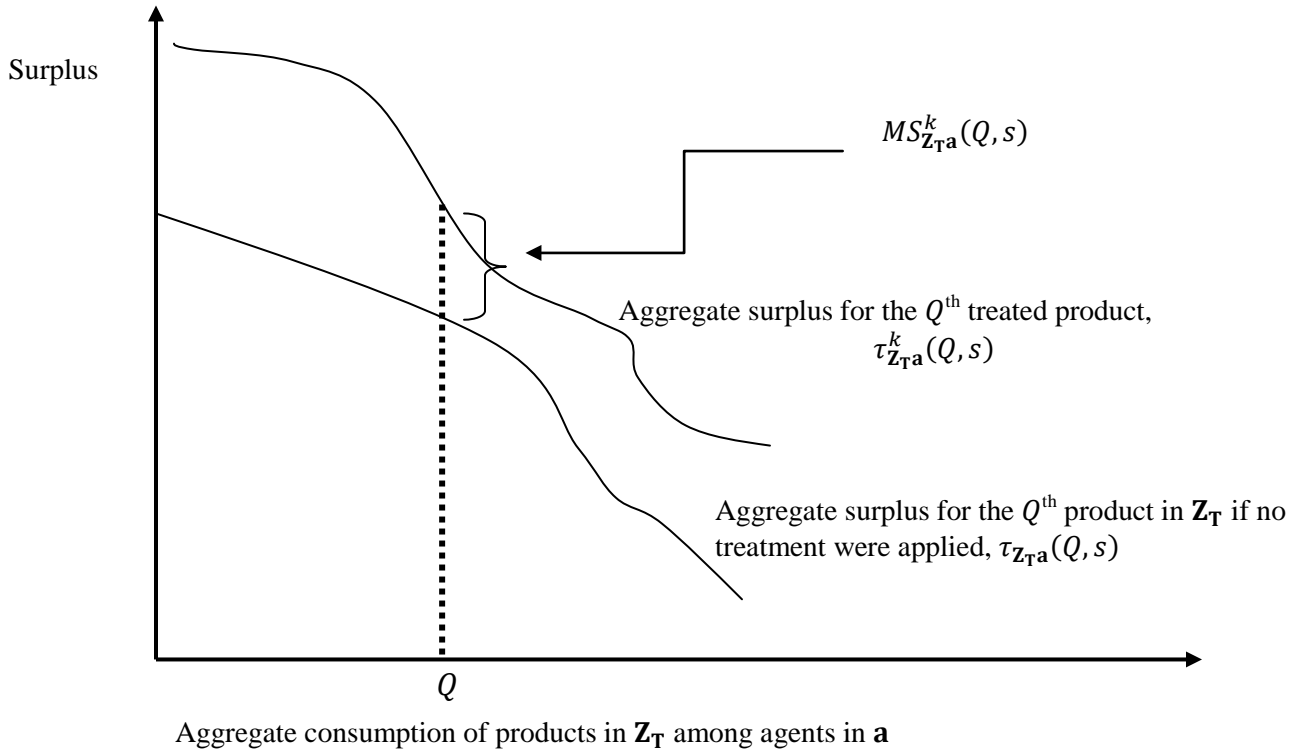
The aggregate benefits of the one-unit increase in z_k to all goods in \mathbf{Z}_T can be expressed in terms of the inverse aggregate treated demand and the inverse aggregate demand. This benefit measure is called the Marginal Surplus (MS):

MARGINAL SURPLUS: The *MS among consumers in \mathbf{a} for a z_k -treatment to goods in \mathbf{Z}_T* is denoted $MS_{\mathbf{Z}_T\mathbf{a}}^k : \mathbb{R}_{[0,+]} \times \mathcal{M}_s \rightarrow \mathbb{R}$. This function $MS_{\mathbf{Z}_T\mathbf{a}}^k(Q, s)$ is defined as $\tau_{\mathbf{Z}_T\mathbf{a}}^k(Q, s) - \tau_{\mathbf{Z}_T\mathbf{a}}(Q, s)$.

For each quantity level Q , $MS_{\mathbf{Z}_T\mathbf{a}}^k$ measures the surplus generated by treating the Q^{th} unit of goods in \mathbf{Z}_T that is demanded by consumers in \mathbf{a} .

Figure 5 plots the two terms in $MS_{\mathbf{Z}_T\mathbf{a}}^k$ as functions of Q . The higher and lower curves depict aggregate demand under the regimes of treatment and no treatment, respectively. At each Q level, the $MS_{\mathbf{Z}_T\mathbf{a}}^k$ is the vertical difference between the two curves. While both curves express willingness to pay among consumers in \mathbf{a} , the identity of the buyer purchasing the Q^{th} unit may change with the application of the treatment. The relationship between the MS and the *bid function* is discussed further in the appendix.

Figure 5: Aggregate Demand for Treated and Untreated Products and the Marginal Surplus



Notes to Figure 5: The lower curve illustrates the inverse aggregate demand for products in \mathbf{Z}_T among consumers in \mathbf{a} as a function of the quantity. At each value Q , this curve returns the surplus obtained from the Q^{th} unit consumed. The higher curve illustrates the same inverse aggregate demand after each product in \mathbf{Z}_T has been “treated” with a one-unit increase in z_k . The difference in the two curves measures the effect of the treatment on the surplus obtained from the Q^{th} unit consumed. In addition to Q , MS and the inverse aggregate demand functions depend on s , which summarizes the state of the economy. MS is different from MWTP because the consumer of the Q^{th} unit in the untreated state may be a different person from the one who consumes the Q^{th} unit in the treated state.

V. IDENTIFICATION OF MWTP DENSITIES

A. An Idealized Experiment

We begin our discussion of identification with an idealized experiment that will be feasible when the researcher can apply the *treatment technology* and charge randomly assigned prices for that technology to different consumers:

IDEALIZED EXPERIMENT. To conduct the *idealized experiment*, the researcher draws a sample of N consumers from the population, where the draws are independent. Each consumer has the option to have the *treatment technology for attribute z_k* applied to every good consumed from \mathbf{Z} . To receive this treatment, the consumer must pay a *treatment price* τ_i , where τ_i is

randomly assigned across consumers. The *treatment price* τ_i has a compact support $[\underline{\tau}, \bar{\tau}]$. F_k^{MWTP} is bounded away from zero and one over this support.

The *idealized experiment* can be applied to measure the MWTP for home improvements, product upgrades, and attributes that are artificially tied to specific houses or jobs, as with school district access or health care coverage.

In the absence of the treatment, i selects the bundle q_{zi}^* of goods in \mathbf{Z} and obtains zero surplus. If the treatment is provided at a price $\tau_i = MTWP_{ki}$, then i is able to purchase the bundle $T_{zk}(q_{zi}^k)$ at a cost of $\theta_i(T_{zk}(q_{zi}^k), p_x^*, w_i^*, u_i^*)$. By the definition of the *bid function* and *consumer surplus*, i receives utility of u_i^* and zero *consumer surplus* at this *treatment price*. At the *treatment price* $\tau_i = MTWP_{ki}$, consumer i is therefore indifferent between selecting and not selecting the treatment. *Local non-satiation* implies that, at any *treatment price* greater than $MTWP_{ki}$, purchasing the treatment would give i negative surplus, and at any *treatment price* less than $MTWP_{ki}$, purchasing the treatment would give i positive surplus.

Identification of $MWTP_{ki}$ in the *idealized experiment* is a straightforward application of a standard nonparametric discrete choice estimator (*cf.*, Pagan and Ullah, 1999, pp. 272-99). Consumer i selects the treatment option if and only if $\tau_i \leq MTWP_{ki}$. At a given treatment price τ , the fraction of consumers who select the treatment option is $Pr_i(MWTP_{ki} \geq \tau)$, where the probability is taken over all consumers i . This probability can be rewritten as $1 - F_k^{MWTP}(\tau)$. Let $Treat_i$ be a dummy for whether i selects the treatment. Let \mathbb{Z}_+ denote the set of positive integers. A consistent kernel estimator for $F_k^{MWTP}(\tau)$ can be constructed following Li and Racine (2007, pp. 182-3, 209-10; 2008):¹⁸

BANDWIDTH AND WEIGHTING KERNEL: A bandwidth $h: \mathbb{Z}_+ \rightarrow \mathbb{R}_+$ is defined as a decreasing function of the sample size N . For simplicity, let the value $h(N)$ be denoted h . This function satisfies the conditions that $\lim_{N \rightarrow \infty} h = 0$ and $\lim_{N \rightarrow \infty} N * h = \infty$. A weighting kernel $\omega: \mathbb{R} \rightarrow \mathbb{R}_{[0,+]}$ is a symmetric, bounded pdf that integrates to one.

ESTIMATOR IN IDEALIZED EXPERIMENT: Given a sample size N , bandwidth h , and weighting kernel ω , for any MWTP value $\tau \in [\underline{\tau}, \bar{\tau}]$, our estimator $\tilde{F}_{kN}^{MWTP}(\tau)$ equals

$$\frac{\sum_{i=1}^N \omega\left(\frac{\tau_i - \tau}{h}\right) * (1 - Treat_i)}{\sum_{i=1}^N \omega\left(\frac{\tau_i - \tau}{h}\right)}.$$

Formulas for the estimators are presented here; proofs of consistency and formulas for the bias, variance, and asymptotic distributions appear in the appendix.

B. Studies of Existing Tradeoffs

¹⁸ This estimator is similar to that used in a variety of discrete choice settings including contingent valuation experiments in which survey participants report how they might act if presented with certain hypothetical tradeoffs (Creel and Loomis, 1997; Crooker and Herriges, 2004; Kristrom, 1990).

In some cases, it is possible to learn about consumers' valuations of an important attribute or public good by examining decisions that specific consumers already face:

EXISTING TRADEOFF EXPERIMENT: Let $\mathbf{Z}_0, \mathbf{Z}_k \subseteq \mathbf{Z}$ be two sets of goods that satisfy $\mathbf{z} \in \mathbf{Z}_0 \Leftrightarrow z_1, \dots, z_k + 1, \dots, z_n \in \mathbf{Z}_k$, so that \mathbf{Z}_k contains the treated variant of every good in \mathbf{Z}_0 . A sample of N consumers is drawn from a population $\mathbf{a} \subseteq \mathbf{I}$ of consumers. The consumers in \mathbf{a} all have utility functions such that their optimal bundles in \mathbf{Z} either include only goods in \mathbf{Z}_0 or only goods in \mathbf{Z}_k . Hence, for all $i \in \mathbf{a}$, at every value of τ , $Q_{\mathbf{Z}_0 i}(\tau, s^*) = Q_{\mathbf{Z}_i}(\tau, s^*)$ or $Q_{\mathbf{Z}_k i}(\tau, s^*) = Q_{\mathbf{Z}_i}(\tau, s^*)$. For each consumer in the sample, prices between the two sets differ by a consumer-specific constant *treatment price* τ_i , so that $p_{\mathbf{z}}(\mathbf{z}) = p_{\mathbf{z}}(z_1, \dots, z_k + 1, \dots, z_n) - \tau_i$ for each $\mathbf{z} \in \mathbf{Z}_0$. There is random or quasi-random variation across consumers in τ_i . The *treatment price* τ_i has a compact support $[\underline{\tau}, \bar{\tau}]$. F_k^{MWTP} is bounded away from zero and one over this support.

In the *existing tradeoff experiment*, the researcher first identifies a population of consumers who face an important binary decision. This decision is modeled as a choice between selecting one's goods in \mathbf{Z} entirely from the set \mathbf{Z}_0 or entirely from the treated variant \mathbf{Z}_k . The researcher selects a random sample of consumers already making this decision and randomly assigns taxes or subsidies for selecting the treated option.

Our estimator in the *existing tradeoff* setting is identical to the estimator used in the *idealized experiment*. The distribution of MWTP is estimated with a nonparametric kernel regression of $1 - Treat_i$ on τ_i . The existing tradeoff estimator identifies $F_k^{MWTP}(\tau | i \in \mathbf{a})$ over the support of τ and consequently generates internally valid estimates of the distribution of MWTP for the subpopulation \mathbf{a} . The degree to which these estimates generalize to the overall population depends on the degree to which \mathbf{a} is a representative subset of \mathbf{I} . If decisions are only observed at one *treatment price* τ , $F_k^{MWTP}(\tau | i \in \mathbf{a})$ can be estimated at the MWTP value τ . If the treatment price is assigned according to a discrete probability distribution, then $F_k^{MWTP}(\tau | i \in \mathbf{a})$ can be estimated at the different τ values assigned in the experiment.

In one example of an *existing tradeoff experiment*, Duflo, *et al.* (2006), consider a population of consumers who must decide whether to invest their tax returns into individual retirement accounts or take them as cash. The authors randomly assign a subsidy for selecting the retirement account, so that the time-deferred asset is subsidized for some consumers and not for others. In a quasi-experimental example, Warner and Pleeter (2001) study the U.S. military's drawdown of personnel in the 1990s. Employees leaving the military could select one of two compensation packages. The first was a lump sum payment, and the second was an annuity. The generosity of the annuity relative to the lump sum payment varied across employees according to a nonlinear function of rank and seniority. In both of these experiments, the set \mathbf{Z}_0 consists of cash today, and the set \mathbf{Z}_k consists of cash in the future. Measuring the value of the time-deferred asset relative to cash today provides estimates of the discount rates for these consumers.

In another example, Abrams and Rohlfs (2011) examine criminal defendants who are deciding whether to post bail or to remain in jail until trial. The market of interest \mathbf{Z} is one's activities over the 90

days until one's trial. Defendants who select the set (and do not post bail) \mathbf{Z}_0 must remain in jail, so that they have relatively little freedom over those 90 days, and defendants who select \mathbf{Z}_k and post bail are allowed to be free from jail over those 90 days. The authors focus on an experiment conducted in 1981 in Philadelphia that generated random variation across defendants in the bail levels that they faced. This exogenous variation in bail makes it possible to measure the value that defendants placed on 90 days of freedom from jail.

An example of a quasi-random discrete probability distribution for the price of treatment appears in a study by Rohlfs, Sullivan, and Kniesner (2013) of consumer valuations of automobile air bags. In that case, the market of interest is new and used vehicles, \mathbf{Z}_0 includes vehicles without air bags, and \mathbf{Z}_k includes the same vehicles with air bags. Due to government regulations, the supply of air bags increased dramatically over a short period, so that in the early periods of the data, the premium on an air bag reflected the valuation of an air bag for someone at the high end of the MWTP distribution. As the air bags became steadily more common, the premium on an air bag reflected the valuation of an air bag for a consumer at the median, and later at the low end of the MWTP distribution, so that the shifts trace out the shape of the MWTP distribution.

C. Offer-Restricted Environments

A related strategy to the existing tradeoffs approach is to induce consumers to participate in an experiment that restricts their choices:

OFFER-RESTRICTED EXPERIMENT: A sample of N consumers is drawn from the population, and each one is offered a large dollar payment δ to participate in the experiment. Each participant must select bundles in \mathbf{Z} that either include only goods in \mathbf{Z}_0 or only goods in \mathbf{Z}_k . Hence, for each i , participation in the experiment requires that $q_{zi}(\mathbf{z}) = 0$ for every $\mathbf{z} \notin \mathbf{Z}_0$ or $q_{zi}(\mathbf{z}) = 0$ for every $\mathbf{z} \notin \mathbf{Z}_k$. The payment δ is sufficiently large that every consumer opts to participate. Each participant may choose between consuming goods in \mathbf{Z}_0 or goods in \mathbf{Z}_k . The researcher randomly assigns a *treatment price* τ_i across participants in the experiment. Let the *price schedule* in the *offer-restricted experiment* be denoted $(p_x^*, p_z^{Z_T})$, where $p_z^{Z_T}(\mathbf{z}) = p_z^*(\mathbf{z})$ for all $\mathbf{z} \in \mathbf{Z}_0$ and $p_z^{Z_T}(\mathbf{z}) = p_z^*(z_1, \dots, z_k - 1, \dots, z_n) + \tau_i$ for all $\mathbf{z} \in \mathbf{Z}_k$. The *treatment price* τ_i has a compact support $[\underline{\tau}, \bar{\tau}]$. The fraction selecting the treated option is bounded away from zero and one over this support.

OFFER-RESTRICTED CONSUMPTION BUNDLE: For a set of goods $\mathbf{Z}_T \subseteq \mathbf{Z}$ and a dollar payment δ , *consumer i 's offer-restricted consumption bundle for a set \mathbf{Z}_T and payment δ* is i 's optimal *consumption bundle* $(q_{xi}^{Z_T}, q_{zi}^{Z_T})$ given the restriction that $q_{zi}^{Z_T}(\mathbf{z}) = 0$ for all $\mathbf{z} \notin \mathbf{Z}_T$ and supposing that i is given a dollar payment δ . Analytically, this consumption bundle is the solution to the following optimization problem: $\max_{q_{xi}, q_{zi}} u_i(q_{xi}, q_{zi})$ subject to $p_x^* \cdot q_{xi} + p_z^{Z_T} \cdot q_{zi} \leq w_i + \delta$ and $q_{zi}(\mathbf{z}) = 0$ for all $\mathbf{z} \notin \mathbf{Z}_T$.

OFFER-RESTRICTED MWTP: Let $\mathbf{Z}_0, \mathbf{Z}_k \subseteq \mathbf{Z}$ be two sets of goods that satisfy $\mathbf{z} \in \mathbf{Z}_0 \Leftrightarrow z_1, \dots, z_k + 1, \dots, z_n \in \mathbf{Z}_k$, so that \mathbf{Z}_k contains the treated variant of every good in \mathbf{Z}_0 . Let δ be a

dollar payment offered to i . *Consumer i 's offer-restricted MWTP for z_k from goods in \mathbf{Z}_0 with payment δ* is scalar-valued and equals $\theta_i(q_{zi}^{\mathbf{Z}_k}, p_x^*, w_i^* + \delta, u_i(q_{xi}^{\mathbf{Z}_0}, q_{zi}^{\mathbf{Z}_0})) - p_z^{\mathbf{Z}_k} \cdot q_{zi}^{\mathbf{Z}_k}$. The pdf $f_k^{\mathbf{Z}_0}: \mathbb{R} \rightarrow \mathbb{R}_{[0,+]}$ denotes the density of the offer-restricted MWTP, and the corresponding CDF is written as $F_k^{\mathbf{Z}_0\delta}: \mathbb{R} \rightarrow [0,1]$.

In the offer-restricted experiment, consumers are invited to join the study and to restrict their choices in \mathbf{Z} to come either entirely from \mathbf{Z}_0 or entirely from \mathbf{Z}_k . In exchange for accepting this restriction, participants are given a payment δ . As with the idealized and existing tradeoff experiments, there is a *treatment price* τ_i that is randomly assigned across consumers, and consumers must pay this price in order to consume goods from \mathbf{Z}_k . Because consumers are paid to participate in the experiment, the offer-restricted MWTP is defined using $u_i(q_{xi}^{\mathbf{Z}_0}, q_{zi}^{\mathbf{Z}_0})$ rather than $u_i(q_{xi}^*, q_{zi}^*)$ as the benchmark utility level.

The estimation strategy in the *offer-restricted experiment* is the same as in the idealized and existing tradeoff experiments. A nonparametric kernel regression of $1 - Treat_i$ on τ_i identifies the CDF of the offer-restricted MWTP over the support $[\underline{\tau}, \bar{\tau}]$. Oftentimes, the researcher will not be able to offer a sufficiently high payment δ to induce every consumer to participate in the study. In such cases, the study will produce internally valid estimates of $F_k^{\mathbf{Z}_0}$ for a selected sample of consumers who are particularly receptive to cash incentives or to the chance to receive the treatment.

One important example of an *offer-restricted experiment* is the Rand Health Insurance Experiment (Manning, *et al.*, 1987). The researchers randomly assigned health insurance plans across participants, so that some consumers faced high prices for doctor and hospital visits and others faced low prices. The authors use this random variation in prices to estimate the willingness to pay for doctor visits and other types of medical care. Another example of an *offer-restricted experiment* is the Internet Demand Experiment (Edell, and Variaya, 1999; Varian, 2001). Consumers participating in that study agreed to have their internet service provided by the researchers. Every time consumers went online, they faced a menu of different amounts of bandwidth, each sold at a different price. These prices for bandwidth were randomly assigned across consumers. A third example of *offer-restricted experiments* involves laboratory or field experiments to measure the discount rate (*e.g.*, Harrison, Lau, and Williams, 2002; McClure, *et al.*, 2004). In such studies, consumers are offered a cash amount to be paid now or a slightly larger amount to be distributed later (*e.g.*, \$100 now or \$100 plus some additional amount in seven months). The additional amount that is offered later is randomly assigned across consumers.

VI. IDENTIFICATION OF AGGREGATE DEMAND FOR ATTRIBUTES

In many cases, no specific consumer faces a choice between treated and untreated sets of goods, but researchers can learn about the demand for treatment by measuring the extent to which it affects total sales of the product of interest. Such estimators can be used to identify MS.

A. Randomized Product Offers

One approach for estimating MS is to generate randomized product offers whose characteristics vary continuously across the consumers being studied:

OFFER DENSITIES. Let the *offer density functions* $g_0: \mathbf{Z} \rightarrow \mathbb{R}_{[0,+]}$ and $g_k: \mathbf{Z} \rightarrow \mathbb{R}_{[0,+]}$ both be pdfs that assign density levels to each good in \mathbf{Z} . These density functions are constructed to satisfy $g_0(\mathbf{z}) = g_k(z_1, \dots, z_k + 1, \dots, z_n)$ for all \mathbf{z} , so that goods drawn from g_k have on average one more unit of z_k than do those drawn from g_0 .

RANDOMIZED OFFER EXPERIMENT. A sample of N consumers is drawn from the population, where N is even. For the first $N/2$ consumers, a good \mathbf{z}_i is randomly selected for each consumer from the distribution g_0 ; for the remaining $N/2$, \mathbf{z}_i is selected according to g_k . Each consumer is offered a subsidy of δ per unit consumed of the offered good, where δ is constant across consumers. Additionally, the researcher randomly assigns a per unit tax τ_i across participants in the experiment, where τ_i has a compact support $[\underline{\tau}, \bar{\tau}]$ with $\bar{\tau} < \delta$. Each good \mathbf{z}_i is offered at a subsidized price of $p_z^*(\mathbf{z}_i) - \delta + \tau_i$. The fraction of consumers who select zero units of the offered good is bounded away from one over $[\underline{\tau}, \bar{\tau}]$ and \mathbf{Z} .

In the randomized offer experiment, each consumer in the sample is offered a different good. In addition to randomly selecting the product offers, subsidies for the offered goods are randomly assigned across consumers.

Let h be a bandwidth and ω be a symmetric weighting kernel. Our parameter of interest and our estimator are defined as follows:

MS FOR THE AVERAGE OFFERED GOOD. The *marginal surplus (MS) for the average offered good at quantity level Q* is denoted $E_{\mathbf{Z}}[MS_{z1}^k(Q, s^*)]$ and equals $\int_{\mathbf{Z}} MS_{z1}^k(Q, s^*) g_0(\mathbf{z}) d\mathbf{z}$.

RANDOMIZED OFFER ESTIMATOR. For any level of Q for which $\int_{\mathbf{Z}} \tau_{z1} \mathbf{1}(Q, s^*) g_0(\mathbf{z}) d\mathbf{z}$ and $\int_{\mathbf{Z}} \tau_{z1}^k \mathbf{1}(Q, s^*) g_0(\mathbf{z}) d\mathbf{z}$ both fall within the interval $[\underline{\tau} - \delta, \bar{\tau} - \delta]$, the estimator $\tilde{E}_{\mathbf{Z}}[MS_{z1}^k(Q, s^*)]$ of MS for the average offered good at quantity level Q equals

$$\text{argmin}_{\tau} \left| Q - \frac{\sum_{i=N/2+1}^N \omega\left(\frac{\tau_i - \tau}{h}\right) * Q_{z_i}(\tau_i - \delta, s^*)}{\sum_{i=N/2+1}^N \omega\left(\frac{\tau_i - \tau}{h}\right)} \right| - \text{argmin}_{\tau} \left| Q - \frac{\sum_{i=1}^{N/2} \omega\left(\frac{\tau_i - \tau}{h}\right) * Q_{z_i}(\tau_i - \delta, s^*)}{\sum_{i=1}^{N/2} \omega\left(\frac{\tau_i - \tau}{h}\right)} \right| + \frac{\sum_{i=N/2+1}^N p_z(\mathbf{z}_i) - \sum_{i=1}^{N/2} p_z(\mathbf{z}_i)}{N/2}.$$

The first argmin term in our random offer estimator estimates the τ value at which consumption of the offered good would equal Q for the average good selected according to the g_k pdf. This argmin measures the *inverse aggregate treated demand* at quantity Q . The second argmin term estimates the τ value at which consumption of the offered good would equal Q for the average good selected according to the g_0

pdf. This argmin measures the *inverse aggregate (untreated) demand* at quantity Q . To estimate these argmins, the researcher first estimates two separate nonparametric kernel regressions of $Q_{z_i}(\tau_i - \delta, s^*)$ on τ_i , one for each of the two halves of the sample. Next, the researcher inverts these functions by conducting a grid search of τ values or using Newton's method. The third term measures the average price difference between the offered goods in the two samples. In the definition of MS, the treatment technology is applied holding prices constant. In this experiment, the prices are different for the average "treated" good drawn from the g_k pdf and the average "untreated" good drawn from the g_0 pdf. The third term in the formula for the estimator corrects for this price difference.

In one recent application of a *randomized offer experiment*, Bertrand, *et al.* (2010) offered subsidized loans to small business owners in South Africa. The researchers randomized multiple features of the loan, including response deadlines, advertising content, and interest rates. Other researchers have implemented *randomized offer experiments* to study the supply of charitable contributions (Karlan and List, 2007; Landry, *et al.*, 2006). In those studies, subjects were approached and solicited for donations; the researchers randomized the physical characteristics of the solicitors and the extent to which the charity offered matching contributions or lottery incentives.

B. Take-up Rate Studies

One common experimental strategy for measuring attribute demands is to examine total sales for matched pairs of products that are similar in all but one attribute. Using such an approach to estimate MS involves conducting two experiments:

TAKE-UP EXPERIMENT: Let $\mathbf{Z}_0, \mathbf{Z}_k \subseteq \mathbf{Z}$ be two sets of goods that satisfy $\mathbf{z} \in \mathbf{Z}_0 \Leftrightarrow z_1, \dots, z_k + 1, \dots, z_n \in \mathbf{Z}_k$, so that \mathbf{Z}_k contains the treated variant of every good in \mathbf{Z}_0 . In a *take-up experiment*, N_{Takeup} consumers are selected from the population. The first $N_{Takeup}/2$ consumers are offered a subsidy of δ per unit of goods consumed from \mathbf{Z}_0 . The second $N_{Takeup}/2$ of consumers are offered a subsidy of δ per unit of goods consumed from \mathbf{Z}_k . Selection into the first or second half of the sample is randomized. Let s^* denote the state of the economy in the *take-up experiment*.

PRICE EXPERIMENT: In the *price experiment*, N_{Price} consumers are selected. Each consumer is charged a tax τ_i , which could be positive or negative, per unit of goods consumed from the set $\mathbf{Z}_\tau \subseteq \mathbf{Z}$. The tax τ_i has a compact support $[\underline{\tau}, \bar{\tau}]$, and the fraction of consumers who purchase zero units of goods from \mathbf{Z}_τ is bounded away from one over this support. Let s^τ denote the state of the economy in the *price experiment*.

There are many different types of studies that have the general structure of the *take-up experiment*, and *take-up experiments* are often easier and less costly to implement than the *idealized, existing tradeoff*, and *offer-restricted experiments* are.

This relative ease of implementation comes at a cost, however. Importantly, it is not possible using the *take-up experiment* to measure the MWTP for any consumers, and we instead focus on estimating *marginal surplus for the marginal unit*:

MARGINAL SURPLUS FOR THE MARGINAL UNIT. For a set \mathbf{Z}_T of goods in \mathbf{Z} , a set \mathbf{a} of consumers, a subsidy of δ per unit of goods consumed in \mathbf{Z}_T , and a state of the economy s , *marginal surplus for z_k for the marginal unit* is defined as $MS_{\mathbf{Z}_T\mathbf{a}}^k(Q_{\mathbf{Z}_T\mathbf{a}}(-\delta, s), s)$.

The marginal surplus for the marginal unit is the vertical difference in aggregate demand between \mathbf{Z}_T and its treated variant, evaluated at the quantity level $Q_{\mathbf{Z}_T\mathbf{a}}(-\delta, s)$ that would be consumed in the absence of the treatment, supposing that all consumers in \mathbf{a} were given a subsidy of δ per unit purchased of goods in \mathbf{Z}_T . This surplus can be viewed as the social benefit associated with applying the *treatment technology* to the last unit of \mathbf{Z}_T that is consumed.

Another limitation with the take-up framework is that, to identify this parameter of interest, it is necessary to impose restrictions on the functional form for *aggregate demand*:

LINEAR DEMAND. *Aggregate demand is linear in goods in \mathbf{Z}_T* if and only if the “own price effect” of the tax τ on quantity demanded of \mathbf{Z}_T is constant with respect to the price. Hence, linear demand for goods in \mathbf{Z}_T implies that, for all s and τ , $\partial Q_{\mathbf{Z}_T\mathbf{a}}(\tau, s)/\partial\tau = \beta_{\mathbf{Z}_T\mathbf{a}}(s)$. The function $\beta_{\mathbf{Z}_T\mathbf{a}}: \mathcal{M}_s \rightarrow \mathbb{R}$ describes the own price effect on quantity demanded as it varies with the state of the economy. This function is constant with respect to τ .

CROSS-SAMPLE VALIDITY OF OWN PRICE EFFECT. For two aggregate demand functions $Q_{\mathbf{Z}_T\mathbf{a}}(\tau, s)$ and $Q_{\mathbf{Z}'_T\mathbf{a}'}(\tau', s')$, with $\int_{\mathbf{Z}_T} 1d\mathbf{z} = \int_{\mathbf{Z}'_T} 1d\mathbf{z}$ and $\int_{\mathbf{a}} 1di = \int_{\mathbf{a}'} 1di$, *cross-sample validity of the own price effect* is satisfied if $\frac{\partial Q_{\mathbf{Z}_T\mathbf{a}}(\tau, s)}{\partial\tau} = \frac{\partial Q_{\mathbf{Z}'_T\mathbf{a}'}(\tau', s')}{\partial\tau}$.¹⁹

For two aggregate demand functions $Q_{\mathbf{Z}_T\mathbf{a}}(\tau, s)$ and $Q_{\mathbf{Z}'_T\mathbf{a}'}(\tau', s')$ to satisfy *cross-sample validity of the own price effect*, the effect of the per unit tax τ on quantity demanded of \mathbf{Z}_T by consumers in \mathbf{a} in state of the economy s at tax level τ must equal the effect of the per unit tax on quantity demanded of \mathbf{Z}'_T by consumers in \mathbf{a}' in state of the economy s' at tax level τ . We require that that $Q_{\mathbf{Z}_0\mathbf{I}}(\tau, s^*)$ is *linear* in τ and that *cross-sample validity of the own price effect* is satisfied between $Q_{\mathbf{Z}_0\mathbf{I}}(\tau, s^*)$ and $Q_{\mathbf{Z}_T\mathbf{I}}(\tau, s^T)$.

¹⁹ We require that the two functions examine the demand for equally-sized sets of goods among equally-sized sets of consumers. Hence, we require that $\int_{\mathbf{Z}_T} 1d\mathbf{z} = \int_{\mathbf{Z}'_T} 1d\mathbf{z}$ and $\int_{\mathbf{a}} 1di = \int_{\mathbf{a}'} 1di$. If this assumption is relaxed, then the estimated own price effect from one economy must be multiplied by a correction factor such as $\left(\frac{\int_{\mathbf{Z}'_T} 1d\mathbf{z} \int_{\mathbf{a}'} 1di}{\int_{\mathbf{Z}_T} 1d\mathbf{z} \int_{\mathbf{a}} 1di}\right)$ in order to accurately estimate the own price effect in the other economy.

Given these two conditions, the own price effect estimated from the *price experiment* is a constant slope that can be applied to the results from the *take-up experiment*:²⁰

ESTIMATED PRICE EFFECT. The *estimated price effect* from the price experiment is the coefficient on τ from an ordinary least squares regression of $Q_{\mathbf{Z}_\tau i}(\tau_i, s^\tau)$ on τ_i . For a given sample size N_{Price} , let this coefficient be denoted $\hat{\beta}_\tau$.

TAKE-UP ESTIMATOR. Our *take-up estimator* of the marginal surplus for the marginal treated unit of \mathbf{Z}_0 is denoted $\widetilde{MS}_{\mathbf{Z}_0 i}^k(Q_{\mathbf{Z}_k i}(-\delta, s^*), s^*)$. This parameter is calculated as

$$-\frac{1}{\hat{\beta}_\tau} \left[\frac{\sum_{i=N_{Takeup}/2+1}^{N_{Takeup}} Q_{\mathbf{Z}_k i}(-\delta, s^*) - \sum_{i=1}^{N_{Takeup}/2} Q_{\mathbf{Z}_0 i}(-\delta, s^*)}{N_{Takeup}/2} \right].$$


The difference inside the brackets is estimated from the *take-up experiment* and measures the effect of the additional unit of z_k on demand for goods in \mathbf{Z}_0 ; however, this effect is measured in quantity units rather than dollars. To obtain a dollar measure, the researcher uses the *price experiment* to measure the effect of exogenous price shocks on quantity demanded. By *cross-sample validity of the own price effect*, $\hat{\beta}_\tau$ is an unbiased estimate of $\beta_{\mathbf{Z}_0 a}(s^*)$, the own price effect for the take-up experiment. Dividing the quantity effect by negative one times the own price effect converts this quantity effect into a dollar-denominated measure of the benefit of the additional unit of z_k .

Figure 6 illustrates this estimation strategy. The lower and higher curves plot demand for goods in \mathbf{Z}_0 and \mathbf{Z}_k . Prices in the *take-up study* are the same for goods in the two sets, and the difference shown along the horizontal axis illustrates the effect of the additional unit of z_k on quantity demanded. Multiplying this difference by negative one times the slope of the demand curve gives MS, the vertical increase in the demand curve. This ratio gives the amount of the dollar subsidy for goods in \mathbf{Z}_0 that would be required to increase quantity demanded by as much as the treatment did.

One straightforward application of *take-up experiments* is audit studies. The classic audit study design involves two job applicants, one black and one white, who are given similar resumes and are trained to act similarly. In this setting, the consumer is the employer, and the product attribute is worker race. The difference in callback or hiring rates between the black and white applicants is used to measure the extent to which employers prefer workers of a certain race, holding other factors constant (Fix and Turner, 1998). Research designs similar to the take-up experiment have also been used to measure banks' valuations of the race and sex of loan applicants and home sellers' valuations of the race and sex of

²⁰ In principle, linearity of demand is not required for identification. Aggregate demand could have curvature of unknown form that could be identified using nonparametric regressions on data from the price experiment. If, however, the price effect is estimated from an entirely different economy from that studied in the *take-up experiment*, it is unlikely that the curvature in demand would be similar across the two experiments. In the special case in which $\mathbf{Z}_0 = \mathbf{Z}_\tau$ and $s^* = s^\tau$, so that a single economy is being studied, the curvature is likely to be similar across the two samples, and relaxing the linearity assumption would be appropriate.

Figure 6: Identification of Marginal Surplus using the Take-up Estimator



potential home buyers (e.g., Hu, *et al.*, 2010; Zhao, Ondrich, and Yinger, 2006). In some studies, applicants are not trained to act similarly, but researchers use propensity score matching to identify pairs of job or credit applicants or home buyers who are similar in attributes other than race or sex. “Correspondence studies” represent one variation on audit studies in which researchers generate resumes with randomized attributes including experience, education, and a signal of the applicant’s race or sexual orientation (Bertrand and Mullainathan, 2004; Weichselbaumer, 2003). In these studies, \mathbf{Z}_0 might represent the set of black workers in the data and \mathbf{Z}_k could represent the set of white workers in the data. In order to convert these hiring effects from quantity units into dollar units, the researcher can use separate estimates of the hiring effect of subsidizing certain workers, as in Woodbury and Spiegelman’s (1987) experimental study that offered bonuses to employers for each worker hired from one of their treatment groups. The resulting estimator identifies the amount of the bonus for hiring a black worker at which the marginal employer would be indifferent between hiring a white or a black worker.

In addition to studying discrimination, the design of the *take-up experiment* can be applied to attrition rates in experiments. An influential study by Philipson and Hedges (1998) suggests that

researchers can use differential attrition rates between treatment and control groups in experiments to measure the extent to which consumers value the treatment. Using a strategy of this form, Rohlfs and Zilora (2013) examine the effect of being in the treatment group in the Tennessee STAR class size experiment on the decision to remain in the study. In that study, children who were assigned to large classes left the public school system at higher rates than did children assigned to small classes. In this context, \mathbf{Z}_0 consists of a single good, education in a large class in Tennessee Public Schools. Similarly, \mathbf{Z}_k consists of a single good, education in a small class in Tennessee Public Schools. The effect of being assigned to a small class on remaining in the school system measures the quantity difference shown on the horizontal axis in Figure 6. The slope of the demand curve is estimated using data from other studies of the effect of receiving a private school voucher on the decision to leave the public school system. Dividing the quantity effect by the voucher effect gives an estimate of the private school voucher amount that would be necessary to cause as many people to leave the school system as being assigned to a large class did. This ratio gives measures the dollar equivalent of the benefit of small classes.

Another potential application of the *take-up estimator* would use data from the Moving to Opportunity Experiment (Kling, Ludwig, and Katz, 2005). The population of interest in that study is individuals living in poor neighborhoods. Consumers in the treatment group were offered Section 8 vouchers (which are long-term rent subsidies) as incentives to move to wealthier neighborhoods. Consumers in the control group faced no such incentive. In this context, \mathbf{Z}_0 represents apartments in wealthy neighborhoods, and \mathbf{Z}_k represents the same apartments supplemented with Section 8 vouchers. The take-up rate in that study measures the extent to which offering the Section 8 vouchers caused consumers to move to wealthier areas. When combined with separate estimates of the slope of the demand for housing in a certain area, the effect of the voucher on moving measures the value that poor consumers placed on Section 8 vouchers. Estimates of the slope of the demand for housing in a specific area can be taken from studies of the migration effects of welfare benefits or state taxes (Gelbach, 2004; McKinnish, 2007; Rohlfs and Thompson, 2013). The resulting estimator measures the increase in welfare benefits or the reduction in taxes that the marginal person moving to a wealthier area would regard as equivalent to a Section 8 voucher.

C. Placebo Consumers

Often, a randomized experiment is not ethical or feasible, and the MS must be identified using a quasi-experiment in which a certain set of goods is treated for every consumer:

AGGREGATE TREATMENT. An *aggregate treatment of attribute z_k to a set of goods \mathbf{Z}_T* involves the application of the treatment technology to goods in \mathbf{Z}_T for every consumer in the economy. This treatment affects prices. The *price schedule* after the application of this treatment is denoted (p_x^k, p_z^k) , and the new state of the economy is denoted s^k . Consumer i 's optimal

consumption bundle after the application of the treatment is denoted (q_{xi}^k, q_{zi}^k) , and i 's new utility level $u_i(q_{xi}^k, q_{zi}^k)$ is denoted u_i^k .

The treatment affects a positive mass of consumers and consequently changes the *price schedule* from (p_x^*, p_z^*) to (p_x^k, p_z^k) . Any two goods may be complements or substitutes, and any of the prices in the economy could change as a result of the *aggregate treatment*. For simplicity, we do not explicitly model the case of endogenous attributes, in which attributes other than z_k change in response to the treatment. Incorporating these endogenous attributes into the model is straightforward and does not alter any of the identification results; however, doing so complicates the notation. The case of endogenous attributes is discussed in the appendix.²¹

One important difficulty that the *aggregate treatment* presents is that the shift in z_k among goods in \mathbf{Z}_T and the resulting shift in the *price schedule* both affect *aggregate demand*. In order to disentangle these two effects, we consider two types of consumers:

AFFECTED AND PLACEBO CONSUMERS. There are two subsets $\mathbf{a}, \mathbf{b} \subset \mathbf{I}$ of consumers, *affected consumers* in the set \mathbf{a} , who value attribute z_k in varying degrees, and *placebo consumers* in the set \mathbf{b} , who place no value on z_k . Among the placebo consumers, given a tax $\tau = 0$ per unit of goods in \mathbf{Z}_T and a state of the economy s^k , aggregate treated demand $Q_{\mathbf{Z}_T \mathbf{b}}^k(0, s^k)$ of goods in \mathbf{Z}_T equals aggregate untreated demand $Q_{\mathbf{Z}_T \mathbf{b}}(0, s^k)$ of those goods.

Given the definition of an *aggregate treatment* and the existence of placebo consumers, the *placebo consumer study* is described as follows:

PLACEBO CONSUMER STUDY. Two economies are observed that are separated by time or geography. The state of the first economy is s^* . The second economy is identical to the first except that an aggregate treatment of z_k is applied to the set of goods \mathbf{Z}_T , so that the state of the economy is s^k . In each economy, there is a set \mathbf{a} of affected consumers and a set \mathbf{b} of placebo consumers. Samples of N_a^* affected and N_b^* placebo consumers are drawn from the first economy, and samples of N_a^k affected and N_b^k placebo consumers are drawn from the second economy.

In order for the placebo consumers' behavior to be informative, it is necessary that their responses to the shift in prices are similar to the affected consumers' responses. Analytically, we impose two fixed effects assumptions:

FIXED CONSUMER AND STATE EFFECTS. At a given per unit tax τ for goods in that set, per capita quantity demanded of goods in \mathbf{Z}_T for a group \mathbf{a} or \mathbf{b} of consumers is additively separable into a group-specific component and a component that is specific to the state of the economy. Let the functions $\alpha_{\mathbf{Z}_T \mathbf{a}}: \mathbb{R} \rightarrow \mathbb{R}_{[0, +]}$ and $\alpha_{\mathbf{Z}_T \mathbf{b}}: \mathbb{R} \rightarrow \mathbb{R}_{[0, +]}$ denote the group-specific components of per

²¹ The case of endogenous product attributes is mathematically equivalent to simultaneous changes in prices and wealth. Consider, for example, that, as by-product of the treatment, the good \mathbf{z} is converted into the good \mathbf{z}' . From the consumer's perspective, the new price $p_z^k(\mathbf{z})$ of the good \mathbf{z} might be very high, so that the good \mathbf{z} has become effectively unavailable. The good \mathbf{z}' , however, has become available on the market at the price $p_z^k(\mathbf{z}')$. From a seller's perspective, a consumer who was previously endowed with units of \mathbf{z} that sold at the price $p_z^*(\mathbf{z})$ is now endowed with units of \mathbf{z}' that sell at the price $p_z^k(\mathbf{z}')$.

capita demand for consumers in **a** and **b**. These components vary with the tax rate τ but are constant with respect to the state of the economy s . For a given state of the economy s , let the function $\alpha_{\mathbf{Z}_T s}: \mathbb{R} \rightarrow \mathbb{R}_{[0,+]}$ denote the state-specific component of per capita demand. This function varies with the tax rate τ but is constant across consumers. For each tax rate τ and state of the economy s , $\frac{Q_{\mathbf{Z}_T \mathbf{a}}(\tau, s)}{\int_{\mathbf{a}} 1 di} = \alpha_{\mathbf{Z}_T \mathbf{a}}(\tau) + \alpha_{\mathbf{Z}_T s}(\tau)$ and $\frac{Q_{\mathbf{Z}_T \mathbf{b}}(\tau, s)}{\int_{\mathbf{b}} 1 di} = \alpha_{\mathbf{Z}_T \mathbf{b}}(\tau) + \alpha_{\mathbf{Z}_T s}(\tau)$.

The *fixed consumer and state effects* assumption requires that a shift in prices changes per capita demand of goods in \mathbf{Z}_T in equal amounts for affected and placebo consumers.

The *placebo consumer study* measures the effects of an additional unit of z_k on the quantity demanded of treated goods. As with the *take-up study*, we convert this quantity difference into dollar units using data from the *price experiment*. For the demand estimates from the price experiment to apply to the placebo consumer study, we require *linear demand* and *cross-sample validity of the own price effect* between demand among consumers in **a** for goods in \mathbf{Z}_T in the state of the economy s^τ and for goods in \mathbf{Z}_T in the state of the economy s^* .

Putting these assumptions together, we obtain a straightforward estimator of the *marginal surplus for the marginal treated unit*:

PLACEBO CONSUMER ESTIMATOR. Using data from the placebo consumer study and the price experiment, it is possible to identify $MS_{\mathbf{Z}_T \mathbf{a}}^k(Q_{\mathbf{Z}_T \mathbf{a}}^k(0, s^k), s^k)$. The estimator for this parameter is denoted $\widetilde{MS}_{\mathbf{Z}_T \mathbf{a}}^k(Q_{\mathbf{Z}_T \mathbf{a}}^k(0, s^k), s^k)$ and equals $-\frac{1}{\beta_\tau} \left[\left(\frac{\sum_{i=1}^{N_a^k} Q_{\mathbf{Z}_T i}^k(0, s^k)}{N_a^k} - \frac{\sum_{i=1}^{N_a^*} Q_{\mathbf{Z}_T i}(0, s^*)}{N_a^*} \right) - \left(\frac{\sum_{i=1}^{N_b^k} Q_{\mathbf{Z}_T i}^k(0, s^k)}{N_b^k} - \frac{\sum_{i=1}^{N_b^*} Q_{\mathbf{Z}_T i}(0, s^*)}{N_b^*} \right) \right]$.

The intuition for the *placebo consumers estimator* is similar to that used for the *take-up estimator*. The effect of z_k on quantity demanded is measured with difference-in-differences. The first difference,

$\frac{\sum_{i=1}^{N_a^k} Q_{\mathbf{Z}_T i}^k(0, s^k)}{N_a^k} - \frac{\sum_{i=1}^{N_a^*} Q_{\mathbf{Z}_T i}(0, s^*)}{N_a^*}$, measures the combined effect on quantity demanded of increasing z_k for goods in \mathbf{Z}_T and changing prices from (p_x^*, p_z^*) to (p_x^k, p_z^k) . Because placebo consumers do not value the

additional unit of z_k , the second difference, $\frac{\sum_{i=1}^{N_b^k} Q_{\mathbf{Z}_T i}^k(0, s^k)}{N_b^k} - \frac{\sum_{i=1}^{N_b^*} Q_{\mathbf{Z}_T i}(0, s^*)}{N_b^*}$, only measures the effect of the

price changes. Subtracting the second difference from the first produces an estimate of the effect of increasing z_k for goods in \mathbf{Z}_T on per capita demand for those goods. Multiplying this quantity effect by

$-\frac{1}{\beta_\tau}$ converts it into dollar units, as with the *take-up estimator*.

This difference-in-differences approach can be used to measure the value of any intervention that affects some consumers and not others, where a purchase or action is required to receive the benefits.

Section VII of this study applies this approach to measure the benefit that the marginal military recruit

placed on college tuition benefits. Policies that affect school quality could be evaluated by applying this approach to the housing market, where the placebo group would include non-parents or parents with children outside of the relevant age ranges for the affected school (*cf.* Caetano, 2009; Clauretje and Neill, 2000). To evaluate changes in a job's health care benefits, researchers could examine the effect of the extra benefits on the fraction working in that job, where the placebo consumers would include workers who are already covered through their spouses' jobs (*cf.*, Dey and Flinn, 2008). The placebo consumers strategy can also be used to evaluate policies affecting one group's rights (*e.g.*, same-sex marriage, "sanctuary cities" that do not enforce immigration laws, bans on Burqa-wearing) by measuring the effect of the law on the decision to live in the area for which the law applies. The affected group would constitute the group whose civil rights were affected by the law, and the placebo group would constitute consumers whose rights remained constant.

Recent work by Rohlfs (2012) uses the placebo consumers approach to obtain two estimates of the cost of the Vietnam draft. The first estimate is based upon the decision to attend school in order to avoid being drafted. For young men who were draft age during non-war years, attending college provided only the benefits of education. For young men who were draft age during the height of the Vietnam War, attending college provided the benefits of education *plus* student draft deferments, which offered protection from the draft. Females, who were immune from the draft, are the placebo group (*cf.*, Card and Lemieux, 2001). Using estimates of the effect of the draft on college enrollment, together with estimates of the effects of tuition subsidies on college attendance from Dynarski (2003), the MS of protection from the draft is estimated for the marginal college student. This MS measures the amount of the tuition subsidy that would be required to increase college attendance by as much as the Vietnam draft did.

The second approach in that study examines the effects of draft lottery numbers on the fraction voluntarily enlisting in the military (*cf.*, Angrist, 1991). For a young man who faced a high lottery-induced risk of being drafted, voluntarily enlisting provided the regular costs and benefits of military service *plus* protection from the particularly unsafe and unpleasant service that he could expect if drafted. This effect on enlistment rates is converted into a dollar estimate of MS for the marginal recruit by using separate estimates of the effect of enlistment bonuses on the fraction enlisting in the military. This MS measures the amount of the enlistment bonus that would be required to increase enlistment by as much as the risk of being drafted did.

D. Combined Price and Quantity Estimator

The last estimator developed here is useful when an exogenous shock to z_k does not substantially affect other attributes or substitute or complementary goods but may have price and quantity effects:

COMBINED PRICE AND QUANTITY STUDY. Two economies are observed that are separated by time or geography. The state of the first economy is s^* . The second economy is identical to the first except that an aggregate treatment of z_k is applied to the set of goods \mathbf{Z}_T , so that the state of the economy is s^k . A sample of N_z^* goods is drawn from \mathbf{Z}_T in the first economy, and a sample of N_z^k goods is drawn from \mathbf{Z}_T from the second economy. These goods are indexed by z , and the researcher observes the prices $p_z(\mathbf{z}_z)$ for each of these goods. Additionally, samples of N_q^* and N_q^k consumers are drawn from the first and second economies. The researcher observes quantity demanded of goods in \mathbf{Z}_T for each of these consumers.

In order to identify the *MS for the marginal treated unit*, we require that prices respond in a limited way to the intervention:

LIMITED PRICE EFFECTS. The aggregate treatment only affects prices of goods in \mathbf{Z}_T , and it affects them by a constant amount. Hence, $p_x^k = p_x^*$, $p_z^k(\mathbf{z}) = p_z^*(\mathbf{z})$ for all $\mathbf{z} \notin \mathbf{Z}_T$, and $p_z^k(\mathbf{z}) = p_z^*(\mathbf{z}) + r_k$ for all $\mathbf{z} \in \mathbf{Z}_T$, where the premium r_k on z_k is constant across \mathbf{z} .

In addition, we require that the researcher have data from the *price experiment*, that *linear demand* holds, and that *cross-sample validity of the own price effect* holds between demand among all consumers for goods in \mathbf{Z}_T when the state of the economy is s^T and in \mathbf{Z}_T when the state of the economy is s^* .

Given these assumptions, the price and quantity changes can be combined to estimate *MS for the marginal treated unit*:

COMBINED PRICE AND QUANTITY ESTIMATOR. Our combined price and quantity

estimator $\widetilde{MS}_{\mathbf{Z}_T}^k(Q_{\mathbf{Z}_T}^k(0, s^k), s^k)$, is computed as $\frac{\sum_{z=1}^{N_z^k} p_z(\mathbf{z}_z)}{N_z^k} - \frac{\sum_{z=1}^{N_z^*} p_z(\mathbf{z}_z)}{N_z^*} - \frac{1}{\beta_\tau} \left[\frac{\sum_{i=1}^{N_q^k} Q_{\mathbf{Z}_T}^k(i, s^k)}{N_q^k} - \frac{\sum_{i=1}^{N_q^*} Q_{\mathbf{Z}_T}^*(i, s^*)}{N_q^*} \right]$.

The intuition for this estimator is similar to the logic presented in Section II about bias due to aggregate quantity changes. If only the price and quantity demanded of the product of interest change, then the vertical increase in demand is equal to the change in price minus the slope of the demand curve times the change in quantity demanded.²²

The *combined price and quantity estimator* can be applied in any quasi-experimental setting in which traditional hedonic or discrete choice methods are used. This estimator corrects the bias associated with ignoring aggregate changes in quantity supplied. This estimator is most applicable to products whose attributes are fixed and do not depend upon other consumers' behaviors. Additionally, the *aggregate treatment* should be sufficiently small that it does not affect the prices of complementary goods. Rohlfs,

²² A similar approach can be used if there is a small number of observed attributes or prices that are affected by the aggregate treatment. In that case, the researcher must measure the effect of the treatment on each of those other factors. The researcher must also have separate "attribute experiments" or "price experiments" that measure the effects of those other factors on demand for goods in \mathbf{Z}_T .

Sullivan, and Kniesner (2013) use this quantity correction procedure as a sensitivity test in a study of consumer valuations of automobile air bags. In that case, an estimate of the price elasticity of demand for automobiles is taken from Sallee's (2011) study of the effects of tax incentives on demand for the Toyota Prius.

VII. APPLICATION TO THE GI BILL

To illustrate the placebo consumer estimation strategy, we consider the economic value that military recruits place on educational benefits for military service. A typical hedonic study of the MWTP for educational benefits would measure the extent to which jobs with educational benefits pay lower wages than jobs without those benefits. Relatively few jobs offer such benefits, and those that do (such as service in the military) have substantially different types of compensation, workplace amenities, job requirements, and populations of workers than do jobs that do not offer such benefits. These correlates of job-specific educational benefits would bias labor market estimates of the MWTP for these benefits. A traditional quasi-experimental hedonic study would examine how wages change in a specific job when the benefits are offered. However, wages and benefits are often set simultaneously by the employer. A firm might offer a benefits package to increase its number of workers, holding wages constant, as the U.S. military has done. Additionally, by offering educational benefits, a firm attracts young applicants who value education. The quantity response and the change in the composition of workers would both generate bias in quasi-experimental hedonic studies.

To obtain unbiased estimates of the MS for educational benefits for the marginal military recruit, we use the *placebo consumers estimator*. We use data from two Rand Corporation experiments. The first, conducted by Fernandez (1983), measures the extent to which changing the package of educational benefits offered to new recruits affected enlistment in the military. The second experiment, conducted by Polich, Dertouzos, and Press (1986), measures the extent to which cash bonuses affected enlistment. In both experiments, the benefits were only offered to "high quality" recruits who had high school diplomas and scored above the national median on the Armed Forces Qualification Test, and the benefits were only available to recruits who signed contracts to work in high need specialties. Additionally, in both cases, the benefits packages were randomly assigned across different areas of the U.S., and recruiters were given advertising money to promote the new benefits packages and enlistment in their regions more generally. Asch and Dertouzos (1994) provide a useful comparison of these two studies.

Table 2 describes the benefits packages that were offered through the two experiments. In both cases, the control packages were offered in all areas in the year prior to the experiment. The typical Army contract was for three years; however, high quality recruits could sign up for two-year contracts. The

typical Navy or Air Force contract was for four years; however, three-year contracts were available. The enlistment bonus experiment did not offer bonuses for the shorter-term contracts. The GI Bill Experiment

Table 2: Benefits Packages in Two Military Recruiting Experiments

Panel A: GI Bill Experiment (Fernandez, 1983)	
Experimental Group	Description of Benefits
Control Program	Two-to-one matching of up to \$2,700 in education investments for all high quality recruits. An Army recruit enlisting in a high need specialty may receive a "kicker" of up to \$6,000 in additional matched funds.
Ultra-VEAP Kicker	Same as control, but Army kickers up to \$12,000.
Noncontributory VEAP	Same as control, but DoD covers individual contribution for high quality recruits in high need specialties.
Tuition/Stipend	\$1,200/year in tuition assistance plus \$300/month stipend (both indexed to cost of attendance) for up to four years for high quality recruits in high need specialties.
Panel B: Enlistment Bonus Experiment (Polich, Dertouzos, and Press, 1986)	
Test Cell A (Control)	\$5,000 bonus for high quality recruits signing four-year contracts in high need specialties. \$1,500 to \$3,000 bonuses available to high quality recruits signing four-year contracts in other needed specialties.
Test Cell B	Same as control, but \$8,000 bonus for high quality recruits signing four-year contracts in high-need specialties.
Test Cell C	Same as Test Cell B plus a \$4,000 bonus for high quality recruits signing three-year contracts in high-need specialties.

Notes to Table 2: Costs expressed in nominal (1981 and 1983) dollars.

provided benefits that increased with term length. For the Army, the maximum benefit levels of the Ultra-VEAP kicker, the Noncontributory VEAP, and Tuition/Stipend policies were \$6,000, \$2,700, and \$6,550 more than for the control program. For the Navy and Air Force, the maximum benefits for the Noncontributory VEAP and the Tuition/Stipend were \$2,700 and \$18,510 more than for the control program.²³

The treated set \mathbf{Z}_T of goods consists of enlistment contracts in high need specialties; hence, the decision typically modeled as labor supply will in this case be viewed as demand for a job. Following the original studies, a semi-log (rather than a linear) functional form is used for the aggregate demand equation. For both experiments, the affected group \mathbf{a} consists of high quality potential recruits, and the placebo group \mathbf{b} consists of low quality potential recruits, who were not eligible for the bonuses or educational benefits.²⁴

Panel A of Table 3 presents estimated effects of the benefits packages on enlistment rates.²⁵ Rows 1 to 3 show the estimated effects of each of the educational benefits packages on enlistment rates. Numbers are presented separately for Army and Air Force recruits.²⁶ Rows 4 and 5 show the enlistment effects of bonus packages for three-year and four-year contracts, presented for all services together. Panel B presents estimates of the marginal recruit's valuations of the three educational benefits packages provided in the GI Bill Experiment. The effect of bonuses for three-year contracts is used in the denominator for the Army; the effect for four-year contracts is used for the Air Force. Each estimate measures the bonus level that would increase recruitment by as much as the educational benefits did.

The Hinkley confidence intervals in panel B are calculated according to the asymptotic distribution of the ratio of two normal random variables, as discussed in the appendix. Our estimates of

²³ Assumes 8% annual growth for the indexed Tuition/Stipend benefits, as in Fernandez (1983). To compute these indexed amounts, personnel are assumed to enlist for the standard term length (three years for Army and four years for Air Force and Navy) and begin school immediately after completing service.

²⁴ Many of the high need specialties were only available for high quality recruits. Consequently, the market-level changes (including the effects of advertising) are measured using the placebo groups' total enlistment rates rather than their enlistment rates in certain specialties.

²⁵ The raw data for the difference-in-differences do not appear in either study. In the GI Bill Experiment, the author presents estimated effects of the treatments on high quality enlistments and reports that, in the specifications with controls for local labor market conditions, the treatments had no effect on low quality enlistments (Fernandez, 1983, pg. 55). In the Enlistment Bonus Experiment, the authors use a regression-based approach that measures the effect of the treatment after controlling for month-to-month variation in the numbers of low-quality recruits. For both experiments, we suppose that the estimation procedures correctly identify the effects of the treatments, holding all prices constant. Additionally, in both cases, following Polich, Dertouzos, and Press (1986), the "composite" effect is presented here. This effect measures the log increase that the treatment generated in overall high quality enlistments, controlling for market-level changes, plus the log increase that the treatment generated in the fraction of new recruits who choose the affected specialties and term lengths. The standard errors for these effects are estimated under the assumption that the coefficient estimates are statistically independent. The percentage changes in Fernandez (1983) are converted into log effects using the delta method.

²⁶ Few Navy specialties were covered by the GI Bill Experiment, and the paper does not present separate estimates for Navy recruits.

Table 3: Calculation of the MS for Educational Assistance Packages

Panel A: Log Increase in Enlistment Among Affected Contracts			
	Policy Comparison	Army	Air Force
1.	Ultra-VEAP Kicker relative to Control Program (only Army)	0.24 (0.04)**	
2.	Noncontributory VEAP relative to Control Program	0.08 (0.05)*	0.18 (0.06)**
3.	Tuition/Stipend relative to Control Program	-0.02 (0.05)	0.21 (0.06)**
All Services			
4.	Test Cell B versus Test Cell A (additional \$3,000 bonus for 4-year contract)		0.46 (0.04)**
5.	Test Cell C versus Test Cell B (\$4,000 bonus for 3-year contract)		0.57 (0.12)**
Panel B: Estimated Value of Educational Assistance Packages			
		Army	Air Force
6.	Ultra-VEAP Kicker relative to Control Program (delta method standard error) [Hinkley 95% confidence interval]	\$1,886 (515.0)** [1,297, 3,435]**	
7.	Noncontributory VEAP relative to Control (delta method standard error) [Hinkley 95% confidence interval]	\$638 (394.3) [-90.22, 1,623]*	\$1,260 (638.0)** [378.6, 2,235]**
8.	Tuition/Stipend relative to Control Program (delta method standard error) [Hinkley 95% confidence interval]	-\$160 (378.8) [-1,00.7, 621.8]	\$1,524 (147.2)** [1,282, 1,880]**

Notes to Table 3: Estimates in panel are “composite” log effects, as in page 45 of Polich, Dertouzos, and Press (1986) that add together account the “expansion” effect of the increase in overall enlistments and the “channeling” effect into certain types of enlistment contracts. The valuation estimates in panel B are constructed by dividing the education policy effect by the bonus effect and dividing by $(90.9/99.6)*4,000$ for Army and $(90.9/99.6)*3,000$ for Air Force to obtain a per 1981 dollar effect. The 3-year contract bonus effect is used for Army and the 4-year contract bonus effect is used for Air Force. Standard error and confidence interval approaches are described in the appendix. ** indicates 95% confidence interval does not include zero, * indicates 90% confidence interval does not include zero.

the marginal Army recruit's valuations of the three benefits packages are \$1,886, \$638, and -\$160. These estimates imply valuations of \$0.314, \$0.236, and -\$0.024 per dollar of maximum benefits. For the marginal Air Force recruit, we estimate valuations of the two packages of \$1,260 and \$1,524, or \$0.467 and \$0.082 per dollar of benefits.²⁷

VIII. CONCLUSION

This study extends the theory of hedonic estimation to incorporate three important aspects of markets for heterogeneous goods. First, many important product attributes are endogenous and change in response to exogenous shocks. Second, many heterogeneous goods have complements and substitutes, and exogenous shocks to the market of interest may affect the markets for those other products. Third, aggregate quantity supplied may change in response to an exogenous shock. For all three reasons, the benefits of an exogenous shock to one product attribute will not entirely be capitalized into the price of that product, and traditional hedonic estimators will produce biased estimates. New experimental and quasi-experimental estimators are developed that avoid these biases and are consistent in this general setting. These new estimators are more demanding in their requirements for data and exogenous sources of variation than traditional hedonic estimators are. A variety of such estimators are developed for settings in which different forms of data are available. One of these estimators is applied to measure military recruits' valuations of subsidies for higher education.

REFERENCES

Abrams, David S. and Chris Rohlfs, 2011. "Optimal bail and the value of freedom: evidence from the Philadelphia Bail Experiment," *Economic Inquiry*, 49(3), pp. 750-70.

Alaska Permanent Fund Corporation, 2009. "The Permanent Fund Dividend." Available at: <http://www.apfc.org/home/Content/alaska/dividendPrgrm.cfm>

Angrist, Joshua D., 1991. "The draft lottery and voluntary enlistment in the Vietnam Era," *Journal of the American Statistical Association*, 86(415), pp. 584-95.

²⁷ These per dollar figures are constructed using estimates by Fernandez (1983) of the discounted present values of the different packages. Assuming the maximum contribution and a 20% nominal discount rate, Fernandez estimates that, relative to the control package, the three benefits packages for Army recruits have discounted present values of \$3,200, \$2,100, and \$4,620 (\$0.53, \$0.78, and \$0.71 per dollar of benefits). For Air Force recruits, Fernandez estimates the two benefits packages to have present values of \$1,900 and \$7,640 (\$0.70 and \$0.41 per dollar of benefits) relative to the control package. For all of the packages, the marginal recruit's valuation is considerably smaller than this present value estimate. Three factors explain this discrepancy: 1) the educational benefits are only useful for recruits who would spend that amount of money on college; 2) the recruits may have been skeptical of the unfamiliar Tuition/Subsidy option, and Fernandez notes that they did not understand the inflation adjustment feature of the package; 3) Warner and Pleeter (2001) estimate that enlisted men's nominal discount rates range from 35% to 54%, considerably higher than the 20% assumed by Fernandez.

- Asch, Beth J. and James N. Dertouzos, 1994.** “Educational benefits versus enlistment bonuses: a comparison of recruiting options,” Rand Monograph Report MR-302-OSD. Santa Monica, CA: Rand Corporation.
- Athey, Susan and Guido Imbens, 2007.** “Discrete choice models with multiple unobserved choice characteristics,” *International Economic Review*, 48(4), pp. 1159-92.
- Aumann, Robert J., 1964.** “Markets with a continuum of traders,” *Econometrica*, 32(1/2), pp. 39-50.
- Bajari, Patrick and C. Lanier Benkard, 2005.** “Demand estimation with heterogeneous consumers and multiple unobserved product characteristics: a hedonic approach.” *Journal of Political Economy*, 113(6), pp. 1239-76.
- Bartik, Timothy, J., 1987.** “The estimation of demand parameters in hedonic models,” *Journal of Political Economy*, 95(1), pp. 81-8.
- Bayer, Patrick, Fernando Ferreira, and Robert McMillan, 2007.** “A unified framework for measuring preferences for schools and neighborhoods,” *Journal of Political Economy*, 115(4), pp. 588-638.
- Berry, Steven T. and Philip A. Haile, 2010.** “Nonparametric identification of multinomial demand models with heterogeneous consumers.” Unpublished manuscript.
- Berry, Steven, James Levinsohn, and Ariel Pakes, 1995.** “Automobile prices in market equilibrium,” *Econometrica*, 63(4), pp. 841-90.
- Bertrand, Marianne, Dean S. Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman, 2010.** “What’s advertising content worth? Evidence from a consumer credit marketing field experiment.” *Quarterly Journal of Economics*, 125(1), pp. 263-305.
- Bertrand, Marianne, and Sendhil Mullainathan, 2004.** “Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination,” *American Economic Review*, 94(4), pp. 991-1013.
- Black, Sandra E., 1999.** “Do better schools matter? Parental valuation of elementary education,” *Quarterly Journal of Economics*, 114(2), pp. 577-99.
- Caetano, Gregorio, 2009.** “Estimation of parental valuation of school quality in the U.S.” Unpublished paper.
- Card, David and Thomas Lemieux, 2001.** “Going to college to avoid the draft: the unintended legacy of the Vietnam War,” *American Economic Review*, 91(2), pp. 97-102.
- Carrington, William J., 1996.** “The Alaskan labor market during the pipeline era,” *Journal of Political Economy*, 104(1), pp. 186-218.
- Chay, Kenneth Y. and Michael Greenstone, 2005.** “Does air quality matter? Evidence from the housing market,” *Journal of Political Economy*, 113(2), pp. 376-424.
- Clauret, Terrence M. and Helen R. Neill, 2000.** “Year round school schedules and residential property values.” *Journal of Real Estate and Finance Economics*, 20(3), pp. 311-22.

- Creel, Michael and Loomis, John, 1997.** "Semi-nonparametric distribution-free dichotomous choice contingent valuation," *Journal of Environmental Economics and Management*, 32(3), pp. 341-58.
- Crooker, John and Joseph A. Herriges, 2004.** "Parametric and semi-nonparametric estimation of willingness-to-pay in a contingent valuation framework," *Environmental and Resource Economics*, 27(4), pp. 451-80.
- Cutler, David, Allison B. Rosen, and Sandeep Vijan, 2006.** "The value of medical spending in the United States: 1960-2000," *New England Journal of Medicine*, 355(9), pp. 920-7
- Dey, Matthew and Christopher J. Flynn, 2008.** "Household search and health insurance coverage," *Journal of Econometrics*, 145(1-2), pp. 43-63.
- Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orszag, and Emmanuel Saez, 2006.** "Saving incentives for low- and middle-income families: evidence from a field experiment with H&R Block," *Quarterly Journal of Economics*, 121(4), pp. 1311-46.
- Dynarski, Susan M., 2003.** "Does aid matter? Measuring the effect of student aid on college attendance and completion," *American Economic Review*, 93(1), pp. 279-88.
- Edell, Richard J. and Pravin P. Varaiya, 1999.** "Demand for internet access: what we learn from the INDEX trial," INDEX Project Report #99-009S.
- Ekeland, Ivar, James J. Heckman, and Lars Nesheim, 2004.** "Identification and estimation of hedonic models," *Journal of Political Economy*, 112(S1) pp. S60-109.
- Epple, Dennis, 1987.** "Hedonic prices and implicit markets: estimating demand and supply functions for differentiated products," *Journal of Political Economy*, 95(1), pp. 59-80.
- Fernandez, Richard L., 1983.** "Enlistment effects and policy implications of the Educational Assistance Test Program," Rand Monograph Report R-2935-MRAL. Santa Monica, CA: Rand Corporation.
- Fix, Michael E. and Margery Austin Turner, 1998.** *A national report card on discrimination in America: the role of testing*. Washington, DC: The Urban Institute.
- Gelbach, Jonah B., 2004.** "Migration, the lifecycle, and state benefits: how low is the bottom?" *Journal of Political Economy*, 112(5), pp. 1091-1130.
- Hamermesh, Daniel, 1999.** "Changing inequality in markets for workplace amenities." *Quarterly Journal of Economics*, 114(4), pp. 1085-1123.
- Harrison, Glenn W., Morten Igel Lau, and Melonie B. Williams, 2002.** "Estimating individual discount rates for Denmark: a field experiment," *American Economic Review*, 92(5), pp. 1606-17.
- Heckman, James J., Rosa L. Matzkin, and Lars Nesheim, 2010.** "Nonparametric identification and estimation of nonadditive hedonic models," *Econometrica*, 78(5), pp. 1569-91.
- Hu, Yue, Long Liu, Jan Ondrich, and John Yinger, 2010.** "Do minority and women entrepreneurs face discrimination in credit markets? Improved estimates using matching methods." Unpublished paper.

- Jones, Larry E, 1984.** “A competitive model of commodity differentiation,” *Econometrica*, 52(2), pp. 507-30.
- Karlan, Dean and John A. List, 2007.** “Does price matter in charitable giving? Evidence from a large-scale natural field experiment,” *American Economic Review*, 97(5), pp. 1774-93.
- Kiesel, Kristin and Sofia Villas-Boas, 2007.** “Got organic milk? Consumer valuations of milk labels after the implementation of the USDA organic seal,” *Journal of Agricultural & Food Industrial Organization*, 5(1), article 4.
- Klaiber, H. Allen and V. Kerry Smith, 2009.** “Evaluating Rubin’s causal model for measuring the capitalization of environmental amenities,” National Bureau of Economic Research (Cambridge, MA) working paper no. 14957.
- Kling, Jeffrey R., Jens Ludwig, and Lawrence F. Katz, 2005.** “Neighborhood effects on crime for female and male youth: evidence from a randomized housing voucher experiment,” *Quarterly Journal of Economics*, 120(1), pp. 87-130.
- Kristrom, Bengt, 1990.** “A non-parametric approach to the estimation of welfare measures in discrete response valuation studies,” *Land Economics*, 66(2), pp. 135-9.
- Kuminoff, Nicolai V., and Jaren C. Pope, 2010a.** “Hedonic equilibria, land value capitalization, and the willingness to pay for public goods,” Unpublished paper.
- Kuminoff, Nicolai V., and Jaren C. Pope, 2010b.** “A note on using quasi-experiments to identify hedonic demand parameters,” Unpublished paper.
- Landefeld, J. Steven, and Bruce T. Grimm, 2000.** “A note on the impact of hedonics and computers on real GDP,” *Survey of Current Business*, 80(12), pp. 17-22.
- Landry, Craig E., Andreas Lange, John A. List, Michael K. Price, and Nicholas G. Rupp, 2006.** “Toward an understanding of the economics of charity: evidence from a field experiment,” *Quarterly Journal of Economics*, 121(2), pp. 747-82.
- Lewbel, Arthur, 2000.** “Semiparametric qualitative response model estimation with unknown heteroskedasticity and instrumental variables,” *Journal of Econometrics*, 97(1), pp. 145-77.
- Li, Qi and Jeffrey S. Racine, 2008.** “Nonparametric estimation of conditional CDF and quantile functions with mixed categorical and continuous data,” *Journal of Business & Economic Statistics*, 26(4), pp. 423-34.
- Li, Qi and Jeffrey Scott Racine, 2007.** *Nonparametric econometrics*. Princeton, NJ: Princeton University Press.
- Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz, 1987.** “Health insurance and the demand for medical care: evidence from a randomized experiment,” *American Economic Review*, 77(3), pp. 251-77.
- McClure, Samuel M., David Laibson, George Loewenstein, and Jonathan D. Cohen, 2004.** “Separate neural systems value immediate and delayed monetary rewards,” *Science*, 306(5695), pp. 503-7.

McKinnish, Terra, 2007. “Welfare-induced migration at state borders: new evidence from micro-data.” *Journal of Public Economics*, 91(3-4), pp. 437-50.

Moulton, Brent R., 2001. “The expanding role of hedonic methods in the official statistics of the United States.” Bureau of Economic Analysis, United States Department of Commerce.
Available at: <http://www.bea.gov/about/pdf/expand3.pdf>

Ostroy, Joseph M., 1984. “On the existence of Walrasian equilibrium in large square economies,” *Journal of Mathematical Economics*, 13(2), pp. 143-63.

Pagan, Adrian, and Aman Ullah, 1999. *Nonparametric econometrics*. Cambridge, United Kingdom: Cambridge University Press.

Parmeter, Christopher F. and Jaren C. Pope, forthcoming. “Quasi-experiments and hedonic property value methods.” In John A. List and Michael K. Price, eds., *Handbook on experimental economics and the environment*. Northampton, MA: Edward Elgar Publishing.

Philipson, Tomas and Larry V. Hedges, 1998. “Subject evaluation in social experiments,” *Econometrica*, 66(2), pp. 381-408

Polich, J. Michael, James N. Dertouzos, and S. James Press, 1986. “The enlistment bonus experiment.” Rand Monograph Report R-3353-FMP. Santa Monica, CA: Rand Corporation.

Raff, Daniel M. G. and Manuel Trajtenberg, 1995. “Quality-adjusted prices for the American automobile industry: 1906-1940,” National Bureau of Economic Research (Cambridge, MA) working paper no. 5035.

Roback, Jennifer, 1982. “Wages, rents, and the quality of life,” *Journal of Political Economy*, 90(6), pp. 1257-78.

Rohlfs, Chris, 2012. “The economic cost of conscription and an upper bound on the value of a statistical life: hedonic estimates from two margins of response to the Vietnam draft,” *Journal of Benefit-Cost Analysis*, 3(3), pp. 1-35.

Rohlfs, Chris, Ryan S. Sullivan, and Tom Kniesner, 2013. “New estimates of the value of a statistical life using air bag regulations as a quasi-experiment.” Unpublished paper.

Rohlfs, Chris and Jeffrey P. Thompson, 2013. “Income tax migration: evidence from Connecticut and New York.” Unpublished paper.

Rohlfs, Chris and Melanie Zilora, 2013. “Estimating parents’ valuations of class size reductions using attrition in the Tennessee STAR Experiment.” Unpublished paper.

Rosen, Sherwin, 1974. “Hedonic prices in implicit markets: product differentiation in pure competition,” *Journal of Political Economy*, 82(1), pp. 34-55.

Sallee, James, 2011. “The surprising incidence of tax credits for the Toyota Prius,” *American Economic Journal: Economic Policy*, 3(2), pp. 189-219.

Sheppard, Stephen, 1999. “Hedonic analysis of housing markets.” in Paul Cheshire and Edwin S. Mills, eds., *Handbook of Regional and Urban Economics*, Vol. 3, pp. 1595-1635. Amsterdam; New York and Oxford: Elsevier Science, North-Holland.

United States. Bureau of the Census. 2009. “Population estimates.” Available at: <http://www.census.gov/popest/states/>

United States. Bureau of Economic Analysis. 2010. “State annual personal income.” Available at: www.bea.gov/regional/spi

United States. Bureau of Labor Statistics. 2009. “Average price data.” Available at: <http://www.bls.gov/data>

United States. Energy Information Administration, 2010. “Annual Alaska field production of crude oil.” Available at: http://www.eia.doe.gov/dnav/pet/hist_xls/MCRFPAK2a.xls

Varian, Hal R., 2001. “The demand for bandwidth: evidence from the INDEX project,” in R.W. Crandall and J.H. Alleman, eds., *Broadband: should we regulate high-speed internet access?* American Enterprise Institute-Brookings Joint Center for Regulatory Studies. Washington, DC: American Enterprise Institute.

Viscusi, W. Kip, 1993. “The value of risks to life and health,” *Journal of Economic Literature*, 31(4), pp. 1912-46.

Viscusi, W. Kip, 1996. “Economic foundations of the current regulatory reform efforts,” *Journal of Economic Perspectives*, 10(3), pp. 119-34.

Warner, John T. and Saul Pleeter, 2001. “The personal discount rate: evidence from military downsizing programs,” *American Economic Review*, 91(1), pp. 33-53.

Weichselbaumer, Doris, 2003. “Sexual orientation discrimination in hiring,” *Labour Economics*, 10(6), pp. 629-42.

Woodbury, Stephen A. and Robert G. Spiegelman, 1987. “Bonuses to workers and employers to reduce unemployment: randomized trials in Illinois,” *American Economic Review*, 77(4), pp. 513-30.

Zhao, Bo, Jan Ondrich, and John Yinger, 2006. “Why do real estate brokers continue to discriminate? Evidence from the 2000 Housing Discrimination Study,” *Journal of Urban Economics*, 59(3), pp. 394-419