



Parametric and Nonparametric Approaches to Price and Tax Reform

Angus Deaton; Serena Ng

Journal of the American Statistical Association, Vol. 93, No. 443 (Sep., 1998), 900-909.

Stable URL:

<http://links.jstor.org/sici?sici=0162-1459%28199809%2993%3A443%3C900%3APANATP%3E2.0.CO%3B2-W>

Journal of the American Statistical Association is currently published by American Statistical Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/astata.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Parametric and Nonparametric Approaches to Price and Tax Reform

Angus DEATON and Serena NG

In many public policy problems, we need to estimate the way in which policy changes affect people's behavior. In the analysis of tax and subsidy reform—the topic of this article—we need to know how tax-induced price changes affect the amounts that people buy of the taxed goods. We present various economic and statistical approaches to obtaining the required estimates. We consider the standard structural methods in economics, where the behavior and welfare of individual agents are captured simultaneously by the specification of utility functions whose parameters are to be estimated. We argue that these methods are less useful than alternatives that directly consider the derivatives of the regression function of average behavior. We consider both parametric and nonparametric estimators of these derivatives in the context of price reform for foods in Pakistan, focussing on the advantages and disadvantages of “average derivative estimation” (ADE). ADE is attractive in principle, because it directly estimates the statistics required for policy analysis. In the practical case considered here, neither technique is a clear winner; each has strengths and weaknesses.

KEY WORDS: Average derivative estimator; Demand systems; Nonparametric; Parametric; Tax Reform; Welfare.

1. POLICY CHANGE AND BEHAVIORAL RESPONSES

Much policy analysis is concerned with estimating people's behavioral response to a change in policy. In the case of tax and subsidy reform—the topic of this article—we need to estimate how consumers respond to the changes in price brought about by the policy; for example, how U.S. consumers will change the amount they smoke in response to a tobacco price increase, or how consumers in Pakistan will respond to a reduced subsidy on a staple such as wheat. A case can be made for changing (say, raising) a price if the benefits of doing so exceed the cost. Estimating cost is relatively straightforward, at least for “small” price changes. The money value of the cost to a consumer of a unit increase in the price of a commodity is the quantity consumed; for example, a 1-cent increase in the price of a pack of cigarettes costs 14 cents a week to someone with a two-pack-a-day habit. The benefits of the tax increase come from the social value of the additional revenue generated by the change. This revenue is usefully divided into two parts, the additional revenue that comes from the change in tax with behavior unchanged and the change in revenue that comes from the change in behavior at existing rates. Estimation of the latter poses the greatest challenge.

A simple economic framework can be used to make these ideas more precise. Suppose that the price of good i is p_i , and that this price contains a tax or subsidy element, so that we can write

$$p_i = v_i + t_i \quad (1)$$

where v_i is the price before tax and t_i is the amount of the tax or subsidy. We suppose that t_i can be varied with v_i

held constant, an assumption that must sometimes be modified but that much simplifies a basic presentation. The costs of increasing t_i are conveniently written as

$$\frac{\partial C}{\partial t_i} = \frac{1}{H} \sum_h \xi^h q_i^h. \quad (2)$$

C denotes a suitable measure of social costs, averaged over the individual agents in the economy, labeled $h = 1, \dots, H$. The derivative of the cost of maintaining a constant standard of living with respect to price i for individual h is the amount consumed, q_i^h , and the aggregate cost is obtained by weighting individuals by the amounts ξ^h , the purpose of which is to allow the possibility that policy makers weight different individuals differently, for example by giving greater weight to the poor, to local constituents, or to residents of a particular region. The benefits of the tax increase are represented by the change in revenue, R , again averaged over the population,

$$\frac{\partial R}{\partial t_i} = \frac{1}{H} \sum_h q_i^h + \frac{1}{H} \sum_h \sum_k t_k \frac{\partial q_k^h}{\partial p_i}. \quad (3)$$

Note that the change in the price of i affects not only the demand for goods, but potentially also the demand for the other goods, $k = 1, \dots, n$, and these effects must also be taken into account in the analysis.

The analysis of tax reform (Atkinson and Stiglitz 1980, Dixit 1975, and Newbery and Stern 1987) is concerned with the comparison of (2) and (3). The ratio of (2) to (3) can be interpreted as the social cost per unit of revenue raised through good i , so that if the ratio is low, then the commodity is a candidate for a price rise, and if it is high, then it is a candidate for a tax decrease or subsidy.

This article is concerned with the estimation of these formulas. Equation (2) is straightforward; the weights ξ^h are matters for policy or political judgments, not for estimation, and the quantities consumed by individual (or household) h can be obtained from household survey data, which must

Angus Deaton is William Church Osborn Professor of Public Affairs and Professor of Economics and International Affairs, Princeton University Princeton, NJ 08544. Serena Ng is Associate Professor of Economics, Boston College Chestnut Hill, MA 02167 (E-mail: serena.ng@bc.edu). The authors are grateful to Bo Honoré, Jim Powell, and Tom Stoker for helpful discussions and also to a referee, an associate editor, and the editor for comments on a previous draft. The second author would like to acknowledge Social Science and Humanities Research Council (Canada), Fonds de la Formation de Chercheurs et l'Aide à la Recherche (Quebec), and the Bradley Foundation for financial support.

also contain for each person the information necessary to evaluate the weights. Given that, estimation of (2) is a routine exercise in survey sampling. The first term in (3) can be estimated from the same source, or from administrative tax records, so that only the last term poses any difficulty.

The standard approach in economics—what we call the “structural” approach—is to specify a utility function for each agent h . This utility function tells us simultaneously how the welfare of h is affected by a price change and how its demand is conditioned by price. Section 2 gives a brief summary of how the utility function can be specified and the parameters estimated. It also identifies the main disadvantages of the approach: that the structural parameters (the parameters of preferences) are difficult to estimate without relying on arbitrary and largely untestable assumptions, and that in any case, these parameters lead only indirectly to the quantities that need to be measured for policy. Section 3 considers an alternative strategy in which utility theory is largely ignored, where we start from a parametric specification of the regression function of average, instead of individual demands. We use this method to calculate results for the case study of this article, the analysis of behavioral responses for tax and subsidy reform in Pakistan. Section 4 takes the analysis a step further, using average derivative estimators to provide direct estimates of the average behavioral responses. We discuss the advantages and disadvantages of these nonparametric estimates compared to the parametric analysis in Section 3.

In a good deal of what follows, we simplify the presentation by focusing on one component of the second term on the right side of (3). This is the contribution of the own-price derivative to the change in revenue,

$$t_i \frac{\partial \bar{q}_i}{\partial p_i} = \frac{1}{H} \sum_h t_i \frac{\partial q_i^h}{\partial p_i}. \quad (4)$$

Although it is sometimes important to consider cross-price derivatives—especially when there are closely related goods that are (or might be) taxed at different rates—(4) often accounts for most of the behavioral change. Note that (4) is typically negative—demand curves slope downward—so that the larger absolutely this term, the smaller the revenue increase from the tax increase and the higher the social cost of raising additional revenue via good i . Other things being equal, goods for which (4) is large and negative are goods where subsidies and taxes should be kept to a minimum. Because the other things—(2) and the first term in (3)—are the same for the different estimators considered in this article, the policy consequences of different estimation strategies are confined to the second term in (3), which will often be well approximated by (4).

2. STRUCTURAL APPROACHES: UTILITY-BASED DEMAND SYSTEMS

Although structural, preference-based analyses are used in a wide range of policy exercises in economics, not only in public finance applications, but also in macro-economics, we illustrate with tax reform in developing countries where much of the literature, including that for Pakistan (see Ah-

mad and Stern 1991), assumes that individual behavior conforms to some variant of the *linear expenditure system* (LES) (Stone 1954). Under this specification, household preferences are given by the cost function (the minimum cost of attaining utility u at prices p),

$$x^h = c^h(u^h, p) = \sum_{k=1}^n \gamma_k^h p_k^h + u^h \prod_{k=1}^n p_k^{\beta_k}, \quad (5)$$

where x^h is household total expenditure, u^h is utility, p_k is the price of good k , and the γ_k^h and β_k ($\sum \beta_k = 1$) are parameters to be estimated, the former taken to be household specific and the latter not. The demand functions associated with (5) are

$$p_i q_i^h = p_i \gamma_i^h + \beta_i \left(x^h - \sum_{k=1}^n \gamma_k^h p_k \right). \quad (6)$$

When (5) is inverted to give utility as a function of prices and total expenditure, it provides an explicit formula for the effects of tax-induced price changes on individual welfare. Equation (6) is supplemented by an expression for γ_i^h ; for example, by

$$\gamma_i^h = \zeta_i' z^h + u_i^h \quad (7)$$

where z^h are vectors of observable household characteristics and the stochastic term u_i^h captures unmodeled heterogeneity in demands. Given household level data on demands, prices, and total expenditures (or aggregate average data on the same), together with a specification of u_i^h (or in the aggregate case, their average; often multivariate normality) equations (6) and (7) can be taken to the data and the parameters estimated.

One immediate issue is the restrictiveness of the LES. The limited number of parameters in the model restricts the behavioral responses beyond what is implied by utility theory. In data-poor environments, as when using time-series data in developing countries, this is an advantage. But in general, such a specification largely begs the question with which we began; if individual households conform to the LES, then no matter what the values of the parameters, the cost-benefit calculations will always approve a reduction in tax rate that is higher than the average and approve an increase that is below the average (see Atkinson 1977 and Deaton 1987). In situations where the data are sufficient, the model can be replaced by one of several utility-based flexible functional forms. One such, which we refer to again later, is the almost ideal demand system of Deaton and Muellbauer (1980), whose demand functions take the form

$$s_i^h \equiv \frac{p_i q_i^h}{x^h} = \alpha_i^h + \beta_i \ln \left(\frac{x^h}{P} \right) + \sum_{k=1}^n \theta_{ik} \ln p_k, \quad (8)$$

where P is an appropriately defined price index of all the prices. Formulations like (8) retain the utility basis of models like the LES, so that we retain the integration between welfare and demand systems, but they are much less restrictive and allow the data to affect the policy conclusions.

A much more serious problem with both (6) and (8) is the difficulty in dealing with the fact that the quantities q_i^h (or

expenditures $p_i q_i^h$) are censored at 0. Households do not report *negative* purchases, but in nearly all survey data, most households report zero purchases of at least some goods; indeed, all goods except basic staple foods have some zeros. This phenomenon is inconsistent with either (6) and (7) or with (8) if u_i^h in (7) or α_i^h in (8) are continuously distributed. Although the general theory of utility maximizing behavior allows for the possibility that some goods are not purchased, its application does not lead to the straightforward functional forms for demands as in (6) and (8). Instead, utility must be maximized with the nonnegativity constraints explicitly imposed. In the simplest case, when $n = 2$ and where one of the goods is always purchased, the demand for the other good can often be handled as a censored linear regression model, (as in Heckman 1974), but with $n > 2$ it is necessary to model the selection of goods—typically a polychotomous choice problem—and to recognize that the functional form of the demand function for each good will differ from one selection (or regime) to another. Structural modeling of this type has proved intractable in even the simplest cases (see Lee and Pitt 1986). Even when $n = 2$, a censored linear regression model is a good deal more difficult to estimate credibly than a linear regression model. The standard maximum likelihood (Tobit) estimator is inconsistent under heteroscedasticity or nonnormality. Although this can be repaired using nonparametric Tobit estimators, [such as Powell’s (1984) censored LAD estimator], implementations of the structural approach beyond the empirically uninteresting two-good case, when possible at all, require the use of strong auxiliary assumptions about the distribution of heterogeneity, assumptions that affect the estimates.

Note that even if it were possible to consistently estimate the structural parameters in the presence of censoring, it still remains to aggregate the individual derivatives up to the average required for the policy analysis. To do so requires keeping track of which regime each consumer is in, averaging within each regime separately and allowing for the effect of price changes on the regime selection.

The structural approach thus presents us with two equally unpalatable alternatives. One approach is to estimate simple linear forms, such as (6) and (8), ignoring the censoring problem altogether, treating zero purchases in the same way as other purchases, or deleting them from the analysis. Either way, the theoretical structure—which is seen as the advantage of the approach—is compromised in the estimation, and the parameter estimates obtained will not be consistent for the preference parameters. Alternatively, we can follow the technically much more demanding route of modeling the regimes, and the behavior within each, so that we preserve the theoretical structure, but at the price of obtaining parameter estimates whose consistency is not robust to the changes in the auxiliary distributional assumptions required to obtain them.

3. PARAMETRIC ESTIMATION OF AVERAGE DERIVATIVES

Consider the estimation of the regression function

$$m_i(x, \mathbf{p}, \mathbf{z}) = E(q_i|x, \mathbf{p}, \mathbf{z}), \tag{9}$$

where \mathbf{v} is a vector of household variables on which we wish to condition and which are held constant (along with x) when taxes and prices are changed in computing the derivatives in (3) and (4). If (9) can be estimated on household-level data, then the derivatives that we need can be calculated from

$$\frac{\partial \bar{q}_i}{\partial p_j} = \int \frac{\partial m_i(x, \mathbf{p}, \mathbf{z})}{\partial p_j} dF(x, \mathbf{z}), \tag{10}$$

which can itself be estimated from the sample average of the calculated derivatives. In the parametric work on Pakistan of Deaton (1997, chap. 5), instead of working with the regression functions of q_i and their derivatives, the author works with the regression functions of the *budget shares*, $s_i = p_i q_i / x$, and, based on (8), adopts the following two-equation parametric form:

$$s_i^{ch} = \alpha_i^0 + \beta_i^0 \ln x^{ch} + \gamma_i^0 \cdot \mathbf{z}^{ch} + \sum_{j=1}^n \theta_{ij} \ln p_j^c + (f_j^c + u_{0i}^{ch}) \tag{11}$$

$$\ln v_i^{ch} = \alpha_i^1 + \beta_i^1 \ln x^{ch} + \gamma_i^1 \cdot \mathbf{z}^{ch} + \sum_{j=1}^n \psi_{ij} \ln p_j^c + u_{1i}^{ch}. \tag{12}$$

Here v_i is the “unit value” of good i , defined as the ratio of reported expenditure to reported quantity (e.g., I spent 12 rupees, and bought 4 kilos, so v_i is 3 rupees per kilo). Although (11) looks similar to (8), it has a different interpretation; it is not a structural model of demand conditional on positive purchases, but rather an approximation to the regression function. As such, and unlike (8), its parameters cannot be interpreted as the parameters of preferences.

By (12), unit values are closely related to prices, which are not directly observed, but because richer households buy higher qualities, unit values rise with x and are possibly influenced by other household characteristics, \mathbf{z} , as in (12). Where these effects are absent, $\ln v_i^{ch}$ and $\ln p_i^c$ are identical, the matrix Ψ in (11) is the identity matrix, and all other terms in (12) are 0. More generally, the matrix Ψ allows for “quality shading” in response to price; consumers may respond to higher price by buying not only less, but also lower quality, so that when prices rise, unit values may rise by less.

The three levels of subscripts and superscripts refer to the good, i , the household, h , and the village (or sample cluster, or PSU) in which it lives, c . The unobservable prices, p_j^c , are assumed to be constant within each cluster, f is an unobservable cluster fixed effect, and the u ’s are error terms that are allowed to be correlated across goods and between the share and unit value terms. Such correlation would be present, for example, if when asked to report expenditures and quantities, respondents calculated the former from quantity and prices, so that response errors in prices will be transmitted to reported expenditures.

Estimation of (11) and (12) is done in two stages. At the first, village effects are swept out by working with deviations from village means. This removes not only the fixed

effects, but also the unobservable prices, so that the β and γ parameters can be consistently estimated. The correlations between the paired residuals in share and unit value equations are calculated and interpreted as indicating the variance and covariance of measurement error in the two equations. At the second stage, the estimated β 's and γ 's are used to subtract out the demographic effects from (11) and (12), and the "purged" budget shares and log unit values are averaged, village by village. The between-village regression of "purged" budget shares on "purged" log unit values gives an estimate of price effects, which is corrected for the measurement error calculated at the first step. As is apparent from (11) and (12), the matrices Θ and Ψ are not separately identified, and in fact, the foregoing procedure yields an estimate of only $\Theta\Psi^{-1}$, so that estimation of Θ requires more information. This is provided by a theory of quality shading developed by Deaton (1988; 1997, chap. 5), which uses a separability restriction on preferences that permits the two matrices to be identified separately.

These are complex operations, and it is worth trying to assess which parts matter in practice and which might usefully be short-circuited. Start with the quality issue. As first noted by Prais and Houthakker (1955), these quality effects really do exist, and there is nearly always a significant positive relationship between unit values and total expenditure. The effects are not very large, however. The β coefficients in (10) are estimated to be .10 for both rice and wheat in rural Pakistan, a little higher for dairy products (.14) and meat (.15), but essentially 0 for edible oils and fats and for sugar. Similar results were obtained for Indonesia, Ivory Coast, and the state of Maharashtra in India (Deaton 1997). Numbers of this size make sense. It is hard to imagine a rich person paying much more than twice as much per kilo than a poor person for any broad aggregate of goods, and if rich people (say the top decile) spend about six times as much as poor people (say the bottom decile)—a useful rule of thumb—then the elasticities will be of the size we have estimated. Given that the income elasticities of quality are so low, it is implausible that price elasticities of quality are high, which implies that the matrix Ψ is close to the identity matrix, and so the final correction of the previous paragraph is not very important.

Even if quality effects can be ignored, measurement error remains a real hazard. In particular, Deaton (1988) suggested that regressing logarithms of quantities on logarithms of unit values is not a good idea. But in (11) and (12), the transformation to budget shares and log unit values seems to remove the worst of the measurement error, at least in the sense that the correlations between the residuals of the first-stage within-village regressions are typically not very large, so that the correction for measurement error at the second stage also has little impact. There is no reason that this has to happen, although it is not implausible that although there are correlated measurement errors between reported quantities and reported unit values, there should be relatively little correlation in the reporting errors of expenditures and unit values. Note also that the second-stage regressions are run using village averages, so that the ef-

fects of measurement error, if not eliminated, are likely to be reduced.

The importance of the village fixed effects is harder to judge. Although it is often the case that sweeping out the village means does not affect the estimates of β 's and γ 's—something that is true in Pakistan—the inclusion of village fixed effects is good practice. Villages often differ a great deal from one to another and frequently are internally homogeneous, so that intravillage correlations are to be expected. Furthermore, the village-specific factors, such as prices, are quite likely to be correlated with included characteristics such as income. Indeed, in earlier work with the Ivory Coast, in several cases inclusion or exclusion of village fixed effects had a marked effect on the first-stage estimates.

Note that there is no explicit attempt here to deal with censoring. Zero purchases are included in these regressions, as they must be for the derivatives to be averaged over all households, whether or not they consume each good. Estimation of (11) and (12) will give the correct policy parameters, provided that they correctly specify the regression functions; unlike the structural case, it is not required that the parameters correspond to any parameters of preferences, and in general they will not. Of course, there is no basis (other than convenience) for the assumption that the regression functions take the linear form in (11) and (12).

Table 1 presents results from estimating (11) and (12) using data on seven important food groups from 9,119 rural households in Pakistan's 1984–1985 Household Income and Expenditure Survey. We focus on the most important behavioral terms: the own-price responses, presented here in elasticity form. The wheat, rice, sugar, and edible oils groups choose themselves, because they are the groups whose prices were most distorted in Pakistan at the time of the survey. The other groups are chosen arbitrarily but conform to the usual decomposition of foods in Pakistani surveys. The second column shows the unrestricted own-price elasticities, and the third column shows the same elasticities but with Slutsky symmetry imposed, both sets of estimates calculated following Deaton (1997, pp. 303–314). These are not very different, but we present them to show the structural approach's ability to produce different estimates depending on how much of utility theory we wish to incorporate into the estimation procedures. The symmetric estimates use more of the theory and in some cases (al-

Table 1. Own-Price Elasticities for Seven Goods, Rural Pakistan 1984–1985

	Budget share	Unrestricted	Symmetric
Wheat	.1276	-.61 (.10)	-.63 (.10)
Rice	.0267	-2.16 (.25)	-2.04 (.22)
Dairy products	.1269	-.89 (.04)	-.90 (.04)
Meat	.0366	-.57 (.18)	-.54 (.18)
Oils and fats	.0414	-.80 (.18)	-.81 (.18)
Sugar	.0293	-.11 (.53)	.09 (.53)
Other food	.1219	-.51 (.10)	-.50 (.09)

NOTE: Deaton (1997). Unrestricted estimates of price elasticities are given in column 2, and elasticities with Slutsky symmetry imposed are reported in column 3.

though not here) permit more precise estimates, or even the estimation of responses that cannot otherwise be measured. As incomes rise, the order of consumption of cereals in Pakistan is wheat and then rice, and this pattern is reflected in the price elasticities, with wheat inelastic and rice (the relative luxury good) much more price elastic. (It is worth noting that similar estimates for the state of Maharashtra in India, where wheat is the relative luxury and rice the more basic good, show a much higher price elasticity for wheat than for rice.) Conditional on the parametric specification, the standard errors are small, except for the two groups oils and fats and sugar, whose prices are to some extent controlled and which therefore display much less spatial price variation.

The price elasticities in Table 1 can be inserted into formulas (2) and (3) to calculate desirable directions of tax reform. In fact, the main recommendations are clear from the results and the formulas without detailed calculation. If we are interested in improving the lot of the poorest households, an increase in the price (reduction of the subsidy) of rice will have relatively low cost according to (2) because the poor do not consume rice, but will have relatively high benefit according to (3) because rice is highly price elastic, and the reduction in the subsidy not only will save resources directly, but also will save additional resources as people reduce their consumption. Thus raising revenue through a reduction in the subsidy on rice is estimated to have a high benefit-to-cost ratio. Because wheat is estimated to be a substitute for rice—the cross-effects were estimated but are not shown here—the beneficial effects of raising the price of rice would be moderated somewhat by the switch to wheat, which also carries a subsidy, but the effect is small relative to the equity and own-price effects.

4. NONPARAMETRIC APPROACHES

The ADE to which we now turn requires as its starting point regression functions in which average demands are conditioned on a list of variables, including prices. In many policy applications, such a specification is the natural one with which to begin. In the Pakistani example of this article, the link with Section 2 can be made by starting from (11) and (12) and simplifying as follows. Suppose that prices are equal to unit values, so that we assume away both measurement error and quality issues. Suppose also that the fixed effects in (11) are uncorrelated with the observable conditioning variables, so that (11) implies

$$E(p_i q_i | x, \mathbf{p}, \mathbf{z}) = x \left(\alpha_i^0 + \beta_i^0 \ln x + \gamma_i^0 \cdot \mathbf{z} + \sum_{j=1}^n \theta_{ij} \ln p_j \right) \quad (13)$$

The nonparametric approach dispenses with the functional form in (13) and estimates its average derivative directly.

4.1 Average Derivative Estimation

To simplify the presentation, and for this section only, we denote the dependent variable by y , which in the example

represents the expenditure on each of the seven goods in turn, and use \mathbf{z} to denote all the conditioning variables. We can then write, for each i in turn,

$$E(p_i q_i | \mathbf{z}) = E(y | \mathbf{z}) = m(\mathbf{z}), \quad (14)$$

and the quantity that we wish to measure is

$$b_j = E_z \left(\frac{\partial E(y | \mathbf{z})}{\partial z_j} \right) = E_z \left(\frac{\partial m(\mathbf{z})}{\partial z_j} \right). \quad (15)$$

The expression (15) is an *average derivative*, and to estimate it, we follow Härdle and Stoker (1989) and Stoker (1991) and use *average derivative estimation* (ADE). This is not the only way of proceeding—alternatively, for example, we could nonparametrically estimate the regression functions themselves, using splines, kernels, or locally-weighted estimators, and then calculate and average the derivatives—but it provides a convenient, elegant, and direct method of estimation.

ADE is sometimes thought of as a semiparametric technique, perhaps because it is often used to estimate index regression models where the conditional expectation of y is given by $\phi(\mathbf{x}'\beta)$ for unknown function ϕ . But in the current context we are not restricted to such models, and the method produces estimates of average derivatives without having to restrict the functional form. Even so, the rate of convergence of the estimators is much faster than is typically the case for fully nonparametric treatments; like OLS, the estimates converge at rate $H^{-.5}$ and not at H^{-2} , as would be the case if we were estimating a regression function or its derivatives. The more rapid convergence here is because we are estimating only the *average* of the derivatives, so that although derivatives at any given point will typically be estimated much more imprecisely, their average can be relatively accurate.

The theory of ADE is straightforwardly developed as follows. Suppose that the joint density function of the conditioning variables \mathbf{z} is $f(\mathbf{z})$. Denote the “scores” by

$$w_j = -\partial \ln f(\mathbf{z}) / \partial z_j, \quad (16)$$

a quantity that in principle could be evaluated at each \mathbf{z}^h in the sample, generating an $H \times K$ matrix, where K is the number of conditioning variables. Consider now the unconditional expectation

$$E(w_j y) = \int \int w_j(\mathbf{z}) y f(y, \mathbf{z}) dy d\mathbf{z} = - \int \int f_j(\mathbf{z}) y f(y | \mathbf{z}) dy d\mathbf{z}, \quad (17)$$

where the last expression comes from substituting (16) into (17) and then splitting the joint density function into the product of a conditional and a marginal. If this last term is integrated by parts, and if we assume that $f(\mathbf{z})$ is 0 on the boundary, then (17) becomes

$$E(w_j y) = \int \int f(\mathbf{z}) y f_j(y | \mathbf{z}) dy d\mathbf{z} = E_z(\partial m(\mathbf{z}) / \partial z_j). \quad (18)$$

Hence if we knew the w 's in (16) and formed the $H \times K$ matrix \mathbf{W} , we could calculate

$$\hat{b}_1 = H^{-1} \mathbf{W}' \mathbf{y}, \quad (19)$$

which by (18) would converge to the vector of average derivatives (15). Note from (18) that $E(w_j z_k)$ is the derivative with respect to z_j of the expectation of z_k conditional on z , which is δ_{jk} , the Kronecker delta, so that the probability limit of $H^{-1} \mathbf{W}' \mathbf{Z}$ is the identity matrix. Consequently, another consistent estimator of the average derivatives is provided by the ‘‘instrumental variable’’ estimator

$$\hat{b}_2 = (\mathbf{W}' \mathbf{Z})^{-1} \mathbf{W}' \mathbf{y}. \quad (20)$$

According to Stoker (1991), (20) is preferred to (19) because in (20) common biases in the denominator and numerator offset one another in a way that does not occur with (19).

To make either of these estimators feasible requires a method of estimating the scores (16). This first-stage estimation is based on a kernel estimate of the joint density $f(\mathbf{z})$ given by

$$\hat{f}(\mathbf{z}) = \frac{1}{\tau^k H} \sum_{h=1}^H K \left(\frac{\mathbf{z} - \mathbf{z}^h}{\tau} \right), \quad (21)$$

where τ is a bandwidth and $K(\cdot)$ is some suitable kernel function. In this article we use the quartic product kernel

$$K(u) = \prod_{i=1}^K \kappa(u_i),$$

$$\kappa(u) = \frac{15}{16} (1 - |u|^2)^2 I(|u| \leq 1), \quad (22)$$

where $I(\cdot)$ is the indicator function that is 1 if the statement in brackets is true and 0 otherwise. The bandwidth τ controls the degree of smoothing. Once a bandwidth is chosen; the logarithmic derivatives of (21), which are the estimates of the scores (15), are calculated by differentiating (21) and using the data to calculate the resulting formula. This is straightforward but time consuming; for each of the $H \times K$ (in the Pakistani example, $9,119 \times 9$) elements of the matrix \mathbf{W} , the evaluation requires a complete pass through the sample of 9,119 points. In practice, we follow the standard recommendations of Silverman (1986, p. 77–78), and first transform and scale the \mathbf{Z} matrix so that it has a unit variance–covariance matrix, after which we transform the calculated scores back to restore the original dimensions. The bandwidth is chosen according to the recommendations of Powell and Stoker (1996); in this case we set τ to be unity for the transformed data.

Although it would be possible to estimate (20) using the estimate of the score matrix as described, problems will arise where the estimated density is small, because evaluation of the logarithmic derivative requires division by the estimated density. To avoid this problem, Powell, Stock, and Stoker (1989) evaluated a density-weighted average derivative estimator, which corresponds to (15) weighted by $f(\mathbf{z})$. But Powell, et al. were interested only in estimating the average derivatives up to scale so that they could weight

with impunity, whereas in the current case, we need the average derivatives themselves. Thus we adopt an alternative approach, ‘‘trimming’’ the data by deleting the 5% of the observations for which the estimated density is smallest. If the cutoff on the density to achieve this 5% is α , then we can define the $H \times H$ diagonal matrix $\mathbf{\Omega}$ by its typical element,

$$\omega_{hh'} = \delta_{hh'} I(\hat{f}(\mathbf{z}^h) \geq \alpha). \quad (23)$$

Then the estimator that we actually use can be thought of as the weighted instrumental variable (IV) estimator,

$$\hat{b}_3 = (\hat{\mathbf{W}}' \mathbf{\Omega} \mathbf{Z})^{-1} (\hat{\mathbf{W}}' \mathbf{\Omega} \mathbf{y}). \quad (24)$$

The estimates of the average derivatives are \sqrt{H} consistent and asymptotically normal, with asymptotic standard errors given in (3.5) and (3.6) of Härdle and Stoker (1989) under the assumption that the observations are independent and identically distributed. Let the estimated ‘‘residuals’’ corresponding to household h for good i be given by

$$\begin{aligned} \hat{r}_{hi} &= \hat{w}_{hi} y_h \zeta_h + H^{-1} \tau^{-k} \\ &\times \sum_{h'=1}^H \left[\tau^{-1} K' \left(\frac{\mathbf{z}_h - \mathbf{z}_{h'}}{\tau} \right) - K \left(\frac{\mathbf{z}_h - \mathbf{z}_{h'}}{\tau} \right) \right] \hat{w}_{h'i} \\ &\times \frac{y_{h'} \zeta_{h'}}{\hat{f}_{h'}} \end{aligned} \quad (25)$$

where $\zeta_h = I(\hat{f}(\mathbf{z}^h) \geq \alpha)$. Then the sample covariance matrix for the average derivative estimates is given by

$$\hat{V}_0 = H^{-1} \sum_{h=1}^H \hat{\mathbf{r}}_h \hat{\mathbf{r}}_h' \zeta_h - \bar{\mathbf{r}} \bar{\mathbf{r}}'. \quad (26)$$

Although this formula will be robust to heteroscedasticity in the residuals, it does not account for the possibly more serious bias to standard errors that comes from ignoring intracluster correlations, a particularly inappropriate omission in a context where the cluster structure is such an important part of the analysis. It is also known that the bias to the standard errors tends to be largest when the regressors vary little within the clusters, as is the case for the prices here (see Kloek 1981 and Pfeifferman and Smith 1985). The standard errors of the OLS estimates can be readily corrected using a generalization of Huber–White procedures, and the results from our application suggest that the problem is non-trivial, with the robust standard errors typically more than twice the size of those reported. To adapt the standard errors of the ADEs to allow for similar effects, we consider a modified sample variance–covariance matrix, defined as

$$\hat{V}_1 = H^{-1} \sum_{h=1}^H \sum_{h'=1}^H (\hat{\mathbf{r}}_h - \bar{\mathbf{r}}) (\hat{\mathbf{r}}_{h'} - \bar{\mathbf{r}})' \zeta_{hh'}, \quad (27)$$

where $\zeta_{hh'} = 1$ if household h and h' belong to the same cluster.

4.2 Ordinary Least Squares as an Alternative Estimator

One immediate question is the relationship between ADE and OLS, with the comparison between the average derivatives in (15), as estimated by \hat{b}_3 , and the coefficients of the OLS regression of y on z . Although OLS estimates average slopes of a sort, these are not generally consistent for the average derivatives in (15). In general, if we place no restrictions on the shape of the regression functions, then, as shown by Stoker (1986), OLS will consistently estimate the average derivatives iff the z vector is multinormally distributed. (Note that (20) is OLS if the scores are proportional to the z 's, which requires normality.) If the regression function is restricted in some way, (e.g., to be linear or quadratic), then the distributional assumptions on the z 's can be relaxed, but if we are unwilling to parameterize the model—which is the main point in the present case—OLS will not generally deliver consistent estimates of the quantities required for the policy analysis.

Figure 1, suggested to us by T. Stoker, shows an admittedly extreme case but one illustrating that OLS and ADE estimation are different things. The true regression line has a sigmoid shape, and the distribution of z is bimodal, so the data lie entirely within the two ellipses. The ADE is the average slope over the top and bottom arms, whereas the OLS slope is as shown and bears no relationship to the ADE.

4.3 Some Practical Considerations

The nonparametric estimation procedure, like the parametric one, requires some way of dealing with those cases where no unit value is recorded because the household in question did not buy the good. In the parametric estimation, price elasticities were estimated from the cross-cluster variation in unit values, effectively using average within-cluster unit values to represent each cluster. In the same spirit, we have filled in the missing values for each household with the average unit value for those households in the cluster who do make purchases. In the few cases where no one in the village records a purchase, we use the average for the

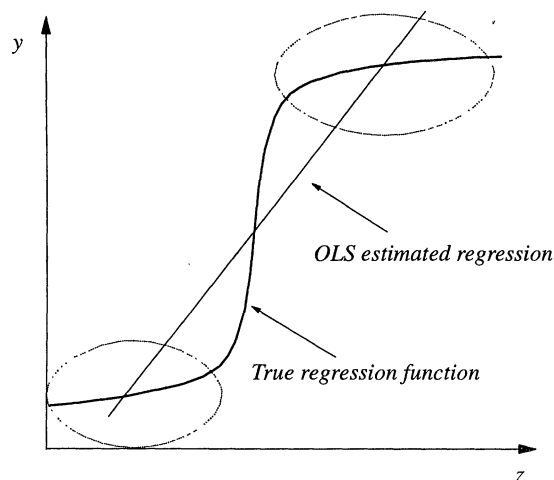


Figure 1. Stoker's Example of the Difference Between OLS and ADE.

province. We must also check the condition for the consistency of ADE that the density of the explanatory variables is 0 at the limits of the regression function. Our variables here are the logarithms of prices and the logarithm of per capita expenditure, which (in logs) appear to be roughly joint normally distributed, but with densities tending to 0 at both high and low values. (Note also that the practical version of the estimator requires trimming of low density points.)

The calculations are straightforward in principle, with no implicit equations to be solved and no iterative techniques required. The computation is burdensome only because of a large number of observations H and a large number of conditioning variables K , and because the number of evaluations is proportional to KH^2 . The baseline calculations, using all observations, took more than one day (over 200,000 seconds) to compute using a 8 processor (MIPS R4400) Unix workstation. This has prompted us to experiment with alternative methods to speed up the computations, all of which involve discarding some of the data. The first method is the simplest and works by sampling a subset of the data randomly. The second and third methods make more deliberate attempts to preserve the structure of the data. Both start by running a preliminary OLS regression of expenditure on the nine conditioning variables, and then ordering the observations by the size of the predicted values. This is done because we need a unidimensional quantity on which to sort the data, but want to avoid selectivity bias induced by selecting on the dependent variable. Method two divides the sorted sample into bins (e.g., quintiles, deciles), then randomly select from each. Method three selects the fifth, tenth, or twentieth observation from the sorted data, starting from some observation other than the first.

Our experience suggests that the computation time is essentially determined by the effective sample size, and not by the method of sample reduction. Estimates from a reduced sample of 2,000 observations took one-third of the time required to obtain estimates from the full sample of 9,119 observations. The point estimates are quite robust to sample reduction. As might be expected, reducing the sample size raises the standard errors of the estimates, so that borderline significant estimates in the full sample tend to become insignificant in the smaller sample. But estimates that are statistically significant in the full sample remain so in the smaller sample. There was no clear winner as to which of the three methods dominate in the sense of giving smaller standard errors. As these sample reduction methods have weak theoretical foundations, we are not particularly confident with these results, which we omit from this article. Suffice it to mention that more sophisticated time-saving methods based on the idea of "discretization" and "convolution" are now being considered in the literature (Härdle and Linton 1994). We remain, however, somewhat skeptical whether such methods are suited for high-dimensional problems as here, because the discretization procedure could itself be very time-consuming.

We have nevertheless found a way of computing the ADE that results in a tenfold reduction in computing time. This comes from replacing the quartic kernel in (23) by the

Gaussian kernel,

$$\kappa(u) = \frac{1}{\sqrt{2\pi}} e^{-u^2/2}. \quad (28)$$

Admittedly, we have not taken extra steps to optimize the execution speed of the programs, as this was not the intent of our article. Indeed, we have been computing the ADEs by “brute force.” Nonetheless, the ADEs can be obtained in a mere 20,000 seconds (instead of 200,000) by replacing the quartic kernel with the Gaussian kernel while leaving the rest of the program unchanged. Our cursory investigation into this difference suggests that the product Gaussian kernel involves one exponentiation and k addition operations. This kernel apparently entails drastically few computer instructions than the quartic kernel, which requires raising variables to higher-power k times. Although not reported in published papers, this time-saving phenomenon has apparently been noted by others (T. Stoker, personal communication). Table 2 reports the full-sample estimates using the Gaussian kernel and the quartic kernel. When evaluated at the same bandwidth of 1.0, the Gaussian estimates are practically identical to the quartic estimates. We report results for the Gaussian kernel for $\tau = .75$ and $\tau = 1.25$, which the readers can verify are sufficiently close to the quartic kernel with $\tau = 1.0$. Clearly, for the application at hand, the estimates are robust to the choice of the kernel and the bandwidth.

4.4 Results and Evaluation

Table 3 shows the bottom line, which is the comparison between the parametric estimates in Section 2 and the ADE of this section, with both sets of results presented as elasticities to show more clearly the different policy implications. Wheat, rice, and sugar are all more price elastic according to the ADE calculations than under the original method. Thus, if these new results are correct, then the subsidies on wheat and rice and the tax on sugar are more distortionary and less desirable than originally estimated.

Table 3 shows the comparison between the OLS regression coefficients for each good and the ADE estimates for two kernels and different bandwidths. The most obvious feature of Table 3 is the close proximity of the OLS coefficients and the ADEs. Although there are a few differences, the two sets of estimates are never very far apart, and when the estimates are converted into elasticities in Table 3 for comparison with the parametric estimates of Section 2, the

only difference of any importance is in the estimate for oils and fats, where the OLS elasticity is -1.31 and the ADE elasticity is -1.77 . The closeness of the OLS and ADE estimates is somewhat disappointing; the object of the article was not to find the most expensive possible way of calculating least squares! As we have seen, this is not a general result. In the current case, the transformation of total expenditure to logs will induce an approximate normality in that variable, but the prices are not obviously normally distributed. Perhaps these results come from a combination of approximate normality for the z 's and approximate linearity for the regression function in this particular application.

The main distinction for the analysis of tax reform in Pakistan is not between OLS and ADE, but between the original parametric method of Section 2 and the non-parametric procedures, taking OLS and ADE together. We currently have no way of deciding which set of estimates should be preferred. The original method deals with issues—measurement error, quality effects, and village fixed effects—that we ignored in the ADEs, whereas the ADEs allow for more general functional forms than permitted in the parametric model.

Consider first the advantages of the ADE method. First, it offers a direct measurement of the quantities that we need. The elasticities in the last column of Table 3 do not come from evaluating the derivatives of some arbitrarily specified functional form at some arbitrarily specified point. Rather, they are estimates of the second term in the denominator of the tax reform formula (3), and although the estimates are interpretable as price elasticities, this is because it is convenient for the purposes of comparison to think of them in such a way, not because we have chosen to parameterize the elasticities. Second, the ADEs do not require that a functional form be specified for the demand functions. We need not concern ourselves with utility theory, nor with the relationship between demand functions and utility functions.

Third, we lose the multitude of problems associated with extensive and intensive margins of consumption. It simply does not matter whether some consumers buy some goods and not others, and there is no need to treat zero demands any differently from positive demands. The welfare and tax reform formulas do not require us to treat zero purchases differently from positive purchases, but depend only on the average consumed and on the derivatives of these averages with respect to price. Because ADEs estimate the average

Table 2. OLS and ADE Estimates of Derivatives of Expenditures With Respect to Log Prices, Rural Pakistan 1984–1985

	OLS	ADE Quartic kernel $\tau = 1.0$	ADE Gaussian kernel $\tau = 1.25$	ADE Gaussian kernel $\tau = .75$
Wheat	11.71 (3.9)	14.33 (5.4) [8.6]	14.39 (5.4) [8.6]	14.16 (5.4) [8.6]
Rice	-46.12 (1.9)	-46.08 (2.6) [5.9]	-46.07 (2.6) [5.9]	-46.08 (2.6) [5.9]
Dairy products	-8.30 (2.4)	-14.97 (3.8) [7.5]	-14.26 (3.9) [7.7]	-14.96 (3.8) [7.5]
Meat	38.30 (2.1)	38.72 (3.0) [5.0]	38.79 (3.0) [5.0]	38.72 (3.0) [5.0]
Oils and fats	-16.66 (7.4)	-42.06 (9.4) [19.6]	-42.10 (9.4) [19.6]	-42.06 (9.4) [19.6]
Sugar	26.76 (8.7)	28.44 (6.7) [13.4]	28.41 (6.7) [13.4]	28.44 (6.7) [13.4]
Other food	79.08 (4.4)	62.77 (6.5) [11.5]	62.77 (6.5) [11.5]	62.77 (6.5) [11.5]

NOTE: OLS is ordinary least squares using all 9,119 observations. ADE is average derivative estimation, τ is the bandwidth. Heteroscedastic-consistent standard errors are in rounded brackets, standard errors corrected for cluster effects are in square brackets.

Table 3. Own-Price Elasticities From OLS and ADE Estimates, Rural Pakistan 1984–1985

	Parametric	OLS	ADE
Wheat	-.61 (.10)	-.93 (.02)	-.91 (.03)
Rice	-2.16 (.25)	-2.31 (.06)	-2.31 (.07)
Dairy products	-.89 (.04)	-1.05 (.01)	-1.09 (.02)
Meat	-.57 (.18)	-.20 (.04)	-.20 (.06)
Oils and fats	-.80 (.18)	-1.31 (.14)	-1.77 (.17)
Sugar	-.11 (.53)	-.30 (.22)	-.26 (.17)
Other food	-.51 (.10)	-.51 (.03)	-.61 (.04)

NOTE: The first column is the second column of Table 1. The second and third columns are calculated from the first and second columns of Table 2.

derivatives directly, and because their consistency does not depend on whether the quantity being affected is zero or positive or has a zero or nonzero derivative, they can be applied directly to quantify the formulas. In the structural approach, we would typically have functional forms for the demand functions and their derivatives, conditional on positive demands, and the calculation of the average derivative would come from integration over the fraction of consumers making positive purchases. (For infinitesimal price changes, nonpurchasers contribute nothing to the derivative, and demand is zero for those whose incomes and characteristics place them on the margin between purchasing or not purchasing at the current price level.) But there is currently no technique available within the utility theory approach that enables us to calculate the set of prices, incomes, and household characteristics that correspond to positive purchases of each good. As a result, the integration is not feasible. On the other hand, the ADE technique renders it unnecessary. Of course, we should be careful about applying ADE estimators to large price changes. For handling such cases, or for going beyond partial equilibrium analysis, there is no alternative to trying to estimate the demand functions themselves.

There are also a number of important disadvantages. The parametric model allows for village fixed effects in demands, thus recognizing that there are likely to be common but unobservable features of behavior in each village, and that these effects may well be correlated with village observables. Even when villages within an area have similar demand patterns, there are likely to be regional or provincial differences that are not simply attributable to differences in income and prices. Although we have not attempted to do so here, the ADE can readily be extended to deal with regional effects, at least in theory. ADEs can be computed for each region and the results averaged using the appropriate regional sampling weights to get an estimate of the national average derivative that we require for the policy analysis. This method could also be extended to allow for the presence of any categorical effect (dummy variable) among the explanatory variables, provided only that the number of observations within each category goes to infinity along with the overall sample size. The average effects of the dummy variables (ADE equivalents for categorical regressors) can then be obtained by calculating an average distance between regression functions for different values of the dummies (see the examples calculated in Deaton and

Paxson 1998). But it remains unclear in the current example whether ADEs can be extended to deal with fixed effects at the level of the village (PSU). Although ADEs can be computed for individual PSUs, and averaged, it is unclear whether such estimates will be consistent if—as is usually the case—PSU size remains fixed as the sample size grows. Estimators can certainly be constructed for special cases, such as additive fixed effects or index models with fixed effects (see Horowitz and Härdle 1994), but these assumptions weaken the nonparametric argument for ADEs.

Second, the parametric estimation method has a procedure for treating the measurement error. Although the effectiveness of the treatment is unclear, and there are examples where the correction makes relatively little difference, the fact remains that within the parametric model there is a range of possibilities for dealing with measurement error, IVs being the most obvious. The use of IVs in semiparametric applications is currently being developed, (see Newey, Powell, and Vella 1995 and Pinkse and Ng 1996), but the use of IVs to deal with measurement error in linear models does not extend in any straightforward way to nonlinear models.

Third, it is unclear how the quality correction procedures in the parametric model can be adapted to the nonparametric case. However, as we have seen, these effects are typically small, and are unlikely to be a major source of difference between the two methodologies. If this were thought not to be the case, ADE estimators could certainly be applied as far as the estimation of the elasticities of quality with respect to total expenditure. It is the next step, where the price elasticities are corrected, that has no obvious counterpart in the nonparametric case.

5. CONCLUSION

The ADEs solve some problems that are intractable in the parametric approach, and although their own problems seem to be addressable, there is a great deal more work to be done. There remain a number of difficulties that are common to both approaches and that cannot be resolved by the nonparametric techniques. In both methods, it is necessary to specify a list of conditioning variables, and the results will typically depend on the choice. Thus there is no possibility of a fully nonparametric treatment, so that the estimation of the behavioral responses in (3) is still on a very different footing from the estimation of the means in (2). Second, the problem of missing unit values is still largely unresolved. In implementing the ADEs, we imputed prices to clusters on the basis of geographical information, a procedure that obviously is sensible and that is supported by the good fit obtained when observed unit values are regressed on cluster dummies. But the parametric approach, by sweeping out cluster fixed effects, requires prices only at the cluster level, so less imputation is required. Furthermore, it seems that different imputations schemes give different results, so that, for example, the OLS estimates in Table 2 are changed nontrivially if regressions are run on cluster means rather than on the individual data. It could also be argued that the need for a parametric imputation

technique for missing values is a good deal less comfortable in a nonparametric setting than in a parametric model. It is also possible that "automatic" imputation techniques could be developed in the nonparametric context. Although there are still very real difficulties, we feel that the use of ADE estimators and semiparametric methods to analyzing tax reforms are sufficiently promising to reward that work.

[Received January 1996. Revised June 1997.]

REFERENCES

- Ahmad, S. E., and Stern, N. H. (1991), *The Theory and Practice of Tax Reform in Developing Countries*, Cambridge, U.K.: Cambridge University Press.
- Atkinson, A. B. (1977), "Optimal Taxation and the Direct Versus Indirect Tax Controversy," *Canadian Journal of Economics*, 10, 590–606.
- Atkinson, A. B., and Stiglitz, J. E. (1980), *Lectures on Public Economics*, New York: McGraw-Hill.
- Deaton, A. S. (1987), "Econometric Issues for Tax Design in Developing Countries," in *The Theory of Taxation for Developing Countries*, eds. D. M. G. Newbery and N. H. Stern, Oxford, U.K.: Oxford University Press, pp. 92–113.
- (1988), "Quality, Quantity, and Spatial Variation of Price," *American Economic Review*, 78, 418–430.
- (1989), "Household Survey Data and Pricing Policies in Developing Countries," *World Bank Economic Review*, 3, 183–210.
- (1997), *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Baltimore: Johns Hopkins University Press.
- Deaton, A. S., and Paxson, C. H. (in press), "Economics of Scale, Household Size, and the Demand for Food," *Journal of Political Economy*, 106.
- Dixit, A. K. (1976), "Welfare Effects of Tax and Price Changes," *Journal of Public Economics*, 4, 103–123.
- Härdle, W., and Linton, O. (1994), "Applied Nonparametric Methods," in *Handbook of Econometrics*, Vol. IV, eds. R. F. Engel and D. F. McFadden Amsterdam: North-Holland.
- Härdle, W., and Stoker, T. M. (1989), "Investigating Smooth Multiple Regression by the Method of Average Derivatives," *Journal of the American Statistical Association*, 84, 986–995.
- Heckman, J. J. (1974), "Shadow Prices, Market Wages, and Labor Supply," *Econometrica*, 42, 679–694.
- Horowitz, J. L., and Härdle, W. (1994), "Direct Semiparametric Estimation of Single-Index Models With Discrete Covariates," mimeo, University of Iowa.
- Kloek, T. (1981), "OLS Estimation in a Model Where a Microvariable Is Explained by Aggregates and Contemporaneous Disturbances Are Equicorrelated," *Econometrica*, 49, 205–207.
- Lee, L.-F., and Pitt, M. M. (1986), "Microeconomic Demand Systems With Binding Non-Negativity Constraints: The Dual Approach," *Econometrica*, 54, 1237–1242.
- Newbery, D. M. G., and Stern, N. H. (1987), *The Theory of Taxation for Developing Countries*, Oxford, U.K.: Oxford University Press.
- Newey, W. K., Powell, J. L., and Vella, F. (1995), "Nonparametric Estimation of Triangular Simultaneous Equations Models," Discussion Paper 95-13, Rice University.
- Pfefferman, D., and Smith, T. M. F. (1985), "Regression Models for Grouped Populations in Cross-Section Surveys," *International Statistical Review*, 53, 37–59.
- Pinkse, J., and Ng, S. (1996), "Nonparametric Two-Step Estimation of Unknown Regression Functions When the Regressors and the Regression Error are not Independent," CRDE Discussion Paper 1596, University of Montreal.
- Powell, J. L. (1984), "Least Absolute Deviations Estimation for the Censored Regression Model," *Journal of Econometrics*, 25, 303–325.
- Powell, J. L., and Stoker, T. M. (1996), "Optimal Bandwidth Choice for Density Weighted Averages," *Journal of Econometrics*, 75, 291–316.
- Powell, J. L., Stock, J. H., and Stoker, T. M. (1989), "Semiparametric Estimation of Index Coefficients," *Econometrica*, 57, 1403–1430.
- Prais, S. J., and Houthakker, H. S. (1955), *The Analysis of Family Budgets*, Cambridge, U.K.: Cambridge University Press.
- Silverman, B. W. (1986), *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall.
- Stoker, T. M. (1986), "Consistent Estimation of Scaled Coefficients," *Econometrica*, 54, 1461–1481.
- (1991), *Lectures on Semi-Parametric Econometrics*, Louvain-la-Neuve, Belgium: CORE Foundation.
- Stone, J. R. N. (1954), "Linear Expenditure Systems and Demand Analysis: An Application to the Pattern of British Demand," *Economic Journal*, 64, 511–527.
- Timmer, C. P. (1981), "Is There 'Curvature' in the Slutsky Matrix?," *Review of Economics and Statistics*, 63, 395–402.
- Timmer, C. P., and Alderman, H. (1979), "Estimating Consumption Parameters for Food-Policy Analysis," *American Journal of Agricultural Economics*, 61, 982–984.