NBER WORKING PAPER SERIES

THE EMPIRICAL ANALYSIS OF TAX REFORMS

Mervyn A. King

Working Paper No. 1996

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 1986

Paper presented to the Symposium on Empirical Public Finance at the Fifth World Congress of the Econometric Socity, MIT, 1985. The research reported here is part of the NBER's research program in Taxation. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

Working Paper #1996
August 1986

The Empirical Analysis of Tax Reforms

ABSTRACT

Over the last decade increasing use has been made of individual household data to analyse the gains and losses from tax reform. Much attention has been paid to the econometric estimation of models of household responses to taxes. But these models yield valid estimates of the welfare consequences of tax changes only when the implied preference orderings are well behaved. This paper discusses the nature of such conditions in detail. Where there are nonlinearities in the budget constraint then two sets of "primal" and "dual" conditions must be satisfied. The analysis of these conditions yields suggestions for the specification of behavioural models and the use of individual-specific information in the observed data.

Professor Mervyn A. King
Department of Economics
London School of Economics
Houghton Street
London, WC2A 2AE
ENGLAND

THE EMPIRICAL ANALYSIS OF TAX REFORMS*

Mervyn A. King

1. Introduction

No subject could be more appropriate or topical for the first World Congress of the Society to be held in the United States than the empirical analysis of tax reform. In May 1985 the President sent his proposals for tax reform to Congress in order "to change our present tax system into a model of fairness, simplicity, efficiency, and compassion, to remove the obstacles to growth and unlock the door to a future of unparalleled innovation and achievement". (US 1985). If enacted, these proposals would make a significant difference to the living standards of many families. The average reduction in taxes as a proportion of income is estimated at 0.6%. But only 58.1% of families would experience a reduction in taxes. (US 1985, Chart 13). It is clear that there are substantial numbers of gainers and losers.

Who gains, who loses? The answer to this question is of interest not only in order to assess the distributional impact of the reform, but also to evaluate why certain proposals are put forward and supported or opposed by different political interest groups. Conventionally, economists assess reforms in terms of the implied change in deadweight loss and typically such estimates are, in terms of order of magnitude, less than one per cent of national income. Figures of this kind are, however, usually an average of large positive and negative values. The mean <u>absolute</u> gain is usually

---

*Financial support was provided by the ESRC Research Programme on Taxation, Incentives, and the Distribution of Income. I am particularly indebted to my colleagues A.B. Atkinson and N.H. Stern for stimulating discussions and comments over many years, and to Ailsa Roell for comments on this paper.

substantially larger and it is the distribution of gains and losses that is relevant to an assessment of the merits of and motives for a reform.

Tax reform can have a major impact on the net incomes and welfare of families, on a par with significant changes in the annual increment to GNP. It is surprising, therefore, that much less attention has been paid to modelling the gains and losses that result from tax reform than has been devoted to macroeconomic modelling. In part this may reflect the acceptance of James Tobin's famous maxim that "it takes a whole heap of Harberger triangles to fill an Okun gap". While there is by definition only one Okun gap there are indeed heaps and heaps of Harberger triangles waiting to be thrown in. Recognition of this has been behind the impetus for tax reform not only in the US but in Europe and much of the English-speaking world. Unfortunately, grandiose claims for the benefits that might be expected to flow from tax reform have been made. How can we assess the plausibility of such claims?

It is obvious, though it can never be stressed too often, that the quality of the model and associated parameter estimates is the critical factor in assessing the plausibility of the simulated gains and losses. A satisfactory model must appear convincing in several dimensions, theoretical consistency, the ability to explain observed empirical phenomena, and a range of econometric criteria (on which there is by no means uniform agreement). This is not simply a question of "flexible functional forms" but of the nature of the model itself. For example, in the study of taxation is it more helpful to work with a static or life-cycle model of labour supply?

In this paper, however, I want to suppose that for the moment we suspend our disbelief in the models that we have estimated, and investigate the consequences of taking them seriously for welfare analysis. I shall argue that this generates additional criteria that should be used when selecting and estimating a model of household behaviour. Two main themes are pursued below. The first is the need for a systematic sensitivity analysis of the consequences of tax changes. The second is the role of individual effects in estimating the distribution of gains and losses as a whole. There is a major difference between time-series analysis on the one hand and the use of cross-section data for welfare analysis on the other. In the latter case our interest is not confined to the model itself, but extends also to the statistical inferences one may draw about the effects of a change in the tax system on each individual in the sample. Individual-specific effects are crucial. It follows that heterogenous preferences are important not only for estimation but also for simulation and welfare analysis. This latter aspect has been neglected relative to the attention that heterogeneity has received in the literature on estimation.

The plan of the paper is as follows. Section 2 examines the role of "nonlinearities" and argues that in this context no representative consumer exists. Aggregate models are, therefore, unhelpful and welfare analysis must be carried out using individual household data. Sections 3-5 discuss the problems of using estimated preference orderings to calculate gains and losses, and illustrate these with reference to two empirical examples in the public finance literature. Finally, section 6 discusses the implications of the issues raised in Section 3 for the estimation of models that are to be used in welfare analysis, and puts forward some tentative suggestions for methods that might be used for both estimation and simulation of tax reforms.

2. Nonlinearities

I shall argue below that some of the problems that arise in the empirical analysis of tax reforms stem from the existence of "nonlinearities" in budget constraints. By nonlinearities I do not mean simply the familiar phenomenon of a nonlinear budget constraint induced by a progressive tax system, but a wide range of factors that affect the opportunity sets of households. Taxes invariably lead to heterogeneity in the prices facing households. In markets in which households can easily transact with each other, of which the best example is asset markets, no equilibrium exists unless nonlinearities of some form are introduced. The government's attempt to price discriminate must be buttressed by tax arbitrage constraints that limit transactions. Examples of this are studied in Miller (1977) and Auerbach and King (1983). Nonlinearities are, therefore, likely to be a generic phenomenon in models of taxation. The main types of nonlinearity are

(i) non-negativity constraints - on, for example, hours worked, consumption demands, asset demands (if short sales are prohibited or assets have different characteristics when held in negative quantities).

(ii) households may not be free to buy or sell as much as they would like at the observed market price. This may be because of explicit rationing or pre-determined contractual arrangements that limit ex post flexibility.

(iii) public goods provision - in which the virtual budget constraint is determined by the marginal willingness to pay for the fixed quantity of the public good.

(iv) discrete choice models - in which the choice between a number of mutually exclusive alternatives cannot be described in terms of a single linear budget constraint.

(v) nonlinearities in the effective budget constraint induced by both the tax schedule and also those benefits paid to households that are contingent upon consumption or labour supply behaviour. The resulting budget sets may be very complicated (Hausman 1981a, 1983b, forthcoming) and Moffitt (1985). Nonconvexities are induced by the interaction between taxes and social security contributions (FICA in the US, NIC in the UK), by special allowances such as the earned income tax credit in the US (Weisbach, 1985), and by programmes for low income families such as AFDC in the US and Family Income Supplement (FIS) and Housing Benefit in the UK.

Nor is the budget constraint necessarily continuous. In Figure 1 we show the budget constraint relating consumption to hours worked for a married man in the UK following the changes to benefits and National Insurance Contributions (NICs) in 1985. The figure is drawn for a married man with a non-working wife, two children, and living in rented accommodation with rent and rates of £15 and £5 a week respectively. Discontinuities are induced by FIS (at 30 hours a week) and NICs. In particular, when the rate of NICs changes at a threshold the new rate is levied on all earnings not just earnings above the threshold. This is true for both employer and employee contributions. Assume that labour market equilibrium equates the marginal product of labour with the total marginal employment cost (the wage rate plus marginal employer's NIC). Then it is natural to assume that the budget constraint should be drawn for a constant marginal product of labour rather than a constant wage rate. With jumps in the NICs rates not only is the budget constraint discontinuous but there are certain ranges of hours worked which are simply infeasible (see Figure 1). Total employment costs jump discontinuously at a NICs threshold and for a given marginal product of labour this implies that the number of hours worked is not a continuous choice variable. Figure 1 is drawn for a marginal product of labour of £2 per hour in 1985 prices.

There are two important consequences of nonlinearities. The first is that there is no representative consumer. Although there are theorems indicating when it is possible to aggregate over endowments (Gorman 1961, Muellbauer 1975, 1976) it is not possible to define a representative consumer for welfare analysis when prices vary across consumers. For example, Figure 1 is by no means representative of the budget constraint facing married men in the UK. The low marginal product of labour was chosen to highlight the discontinuities. This means that both for estimation and welfare analysis we must employ disaggregated household data. In turn, the heterogeneity of household demands and preferences raises a number of issues that are discussed below.

Secondly, with linear budget constraints welfare analysis can proceed only if the conditions for a consistent preference ordering are satisfied by the estimated parameters over the relevant region of the price-income space. But these conditions are not sufficient for welfare analysis to be performed when nonlinearities exist. Nonlinearities may occur at points in the quantity space that are not spanned by the estimated preference ordering. A complete preference ordering may not be recoverable. This is discussed further below.

3. Estimation of Gains and Losses

The result of a change in the tax system can be described in terms of the vector g of gains (or losses) of each household in the economy. In this section we consider some conceptual problems that arise when drawing inferences about g from econometric estimates of household behaviour.

Much of the empirical literature on the distribution of the tax burden has been concerned with the effects of taxes on the net cash incomes of households (Musgrave et al. 1974, Pechman and Okner, 1974, Pechman, 1985). Attention has also been paid, particularly in the analysis of labour supply, to the value of "full income" the market value of a household's total endowment including its time endowment. These measures are valuable for two reasons. First, they focus attention on the diversity of household experience in contrast to the aggregate measures of deadweight loss. Secondly, they are robust with respect to assumptions about household behaviour. Nevertheless, neither cash nor full income are satisfactory measures of gains and losses for reasons which are common to both. A household's welfare, and hence the gain which it experiences following a tax reform, depends upon both its endowment and the prices that it faces. Cash income ignores the latter, and full income incorporates price effects only to the extent that they affect the value of a household's total endowment. Consider, for example, two households identical in all respects except for the wage rate which they can earn. The household with the higher wage rate has a higher full income. But this exaggerates the difference between the two households because, of course, the same household faces a higher price for leisure.

When preferences are defined over a vector of commodities rather than a scalar such as income (so that relative prices matter) then a measure of welfare is required. For empirical purposes it is convenient if welfare can be measured in units that have some natural interpretation, such as dollars or the number of baskets of a given consumption bundle. These correspond to money metric utility (McKenzie 1956, Samuelson 1974) and quantity metric utility (Debreu 1951, 1954, Diewert 1981) respectively. The choice between the two is a matter of taste. In what follows I shall work principally with the dual version of money metric utility proposed by King (1983a, equivalent income function) and Varian (1984, indirect income compensation function). The differences of terminology are of less importance than the fact that both measures are defined over the observable variables that characterise the household's opportunity set. We make the following assumptions.

A1 Households have identical preferences described by the direct and indirect utility functions

$$u = u(\underline{x}) \tag{3.1}$$

$$v = v(\underline{p}, y) \tag{3.2}$$

where $\underline{x}$ denotes the vector of commodity demands (including leisure), $\underline{p}$ denotes the vector of consumer prices (including the wage rate) and y is exogenous full income. Variations in preferences, which have played an important role in the econometric literature, are discussed below.

A2 There exist observable budget constraints for all households. Where the budget constraint is complex it may be difficult in practice to observe its shape (Heckman 1983). For budget constraints that are either nonlinear, nonconvex or discontinuous, it is possible to construct the equivalent continuous linear budget constraint. Examples are given below. This virtual budget constraint is defined by the values of (strictly positive) virtual prices (Rothbarth 1941) and virtual income (Hausman 1981a, 1981b) such that at these virtual prices and income the household would choose the same consumption bundle as it selects under the nonlinear budget constraint. We shall assume for the moment that such a virtual budget constraint exists for all households. But, as we show below, this is not necessarily true and the consequences of nonexistence for the welfare analysis of tax reforms are nontrivial.

To compare welfare levels under different consumption possibility sets we choose a reference price vector, denoted by $\underline{p}^R$. Although the choice of a reference price vector is arbitrary, certain choices, such as current prices, allow a more natural interpretation of gains and losses than others.

Indirect money metric utility, or equivalent income, y_E, as it will be called here, is defined as that level of income which, at the reference price vector, affords the same level of utility as can be attained under the budget constraint ($\underline{p}$, y).[1] Formally,

$$v(\underline{p}^R, y_E) = v(\underline{p}, y) \tag{3.3}$$

Inverting the indirect utility function we obtain equivalent income in terms of the expenditure function

$$y_E = e(\underline{p}^R, v) \tag{3.4}$$

Substituting from the indirect utility function gives equivalent income as a function of reference prices and actual prices and income

$$y_E = f(\underline{p}^R, \underline{p}, y) \tag{3.5}$$

The principal advantage of working with the equivalent income function (EIF) is that it allows a separation of preferences, which are characterised by the functional form of f, and opportunities, which are described by the arguments of f. It is also measured in money units. In the study of income distribution and poverty, attention has traditionally been focussed on measures of resources rather than welfare. To some extent this dichotomy is rather artificial. Equivalent income is itself a measure of resources but computed at a standardised gradient of the budget set (the reference price vector). In principle, it allows us to take into account the various types of "nonlinearities" identified in section 2, but this is feasible only when there exists a virtual budget constraint. The use of the EIF as a measure of resources, however, depends upon the acceptance of individual preferences as the appropriate basis for analysis. Where this is not the case alternative measures may appear more attractive (Atkinson 1985).

The EIF may be used to analyse either the optimal design of the tax system or the consequences of piecemeal reform.[2] From the empirical point of view the latter is usually of more interest. Consider a tax reform that maps the original budget set $(\underline{p}^o, y^o)$ into the post-reform budget set $(\underline{p}^P, y^P)$.[3] A measure of the welfare gain resulting from the reform is given by

$$WG = f(\underline{p}^R, \underline{p}^P, y^P) - f(\underline{p}^R, \underline{p}^0, y^0) \qquad (3.6)$$

Consider two possible states of the economy, s_1 and s_2. The welfare gain to the household in moving from s_1 to s_2 is denoted by $WG_{s_1 \rightarrow s_2}$. There is an infinite number of such measures, each one corresponding to a different reference price vector. Some are more familiar than others. For example, if we take initial prices $\underline{p}^{\cdot}$ as the reference price vector then WG is similar to a Hicksian equivalent variation measure that incorporates changes in income. With post-reform prices $\underline{p}^P$ as the reference price vector WG is the Hicksian compensating variation augmented to allow for income changes. The properties of WG that are important if it is to play a useful role in the empirical analysis of tax changes are

a) For every possible tax reform $s_i \rightarrow s_j$, if utility is higher in state j than in state i then

$$WG_{s_i \rightarrow s_j} > 0$$

This follows directly from the fact that there is, for a given $\underline{p}^R$, a one-to-one correspondence between f and utility.

b) The welfare gain in moving from state i to state j is equal in magnitude though opposite in sign to that in moving in the reverse direction.

$$WG_{s_i \rightarrow s_j} = -WG_{s_j \rightarrow s_i}$$

c) Transitivity : if $WG_{s_i \rightarrow s_j}$ and $WG_{s_j \rightarrow s_k}$ are both positive then not only is $WG_{s_i \rightarrow s_k}$ positive also but the following linear relationship holds

$$WG_{s_i \rightarrow s_k} = WG_{s_i \rightarrow s_j} + WG_{s_j \rightarrow s_k}$$

This can be seen by direct substitution from (3.6).

It should be noted that properties (b) and (c) are <u>not</u> satisfied by the Hicksian equivalent and compensating variations. These measures implicitly assume a different reference price vector for each pairwise comparison of states. For example, the equivalent variation employs the prices in state i to compute the gain in moving from state i to either state j or state k, but the prices in state j to compute the gain in moving from state j to state k. As a result transitivity will not in general hold.[4] It is important, therefore, that once a reference price vector has been chosen, it should remain constant throughout the analysis and also, if possible, when comparing the results of different empirical studies. My own empirical experience (King 1983a, 1983b) has been that estimates of welfare gains are rather insensitive to the choice of reference price vector within the range defined by the pre- and post-reform values, <u>provided</u> that all allocations (both initial state and all possible reforms) are evaluated at the same reference price vector.

If the reform is revenue-neutral then in an economy of identical households WG measures the change in deadweight loss. With heterogeneous households the aggregate change in deadweight loss may be defined as the sum of welfare gains. Even if households have identical preferences and face a common set of prices (including wage rates), this aggregate measure is independent of the distribution of income only when preferences are of the Gorman (1961) polar form (parallel linear Engel curves). In general, therefore ranking reforms in terms of their effect on deadweight loss implies a set of distributional weights. Focussing on per capita deadweight loss (the mean value of WG) ignores most of the information contained in the vector of welfare gains. This contrasts sharply with the sophisticated treatment of individual-specific differences that characterises recent econometric study of micro-data (Heckman 1974, Hausman and Taylor 1981).

Since the equivalent income (or indirect compensation) function is simply an expenditure function evaluated at reference prices and defined explicitly over the actual budget constraint it must satisfy certain properties if it is to represent a consistent preference ordering. These are

(i) f is homogeneous of degree one in $\underline{p}^R$

(ii) f is homogeneous of degree zero in $\underline{p}$ and y

(iii) f is nondecreasing in the reference prices $\underline{p}^R$.

(iv) f is nonincreasing in the prices $\underline{p}$.

(v) f is increasing in y

(vi) f is concave in reference prices

(vii) f is continuous with first and second derivatives (except possibly on a set of measure zero)

(viii) f satisfies the boundary condition

$$f(\underline{p}, \underline{p}, y) = y$$

(ix) Demands may be obtained from f in two ways:

a) Given that f is an expenditure function defined over $\underline{p}^R$ and v, then from Shephard's lemma

$$x_j = \frac{\partial f}{\partial p_j^R}\bigg|_{\underline{p}^R=\underline{p}} \qquad (3.7)$$

b) Differentiating f with respect to p_j holding v and $\underline{p}^R$ constant yields

$$x_j = \frac{-\partial f/\partial p_j}{\partial f/\partial y} \qquad (3.8)$$

If conditions (i) to (viii) are satisfied then the demand functions (3.7 or 3.8) may be integrated to yield a functional form for the EIF. (Hurwicz and Uzawa 1971). It is difficult, however, to find functional forms that satisfy these conditions either globally or over the feasible price-income space. So-called "flexible functional forms" satisfy the conditions only over certain regions of the price-income space which in turn depend upon the estimated parameters. Some of the properties are, however, easily imposed and arise naturally from the derivation of the EIF from an estimated demand system. These are the homogeneity and monotonicity properties (i), (ii), (vii), (viii) and (ix). We shall assume that in empirical work these properties are imposed. The remaining properties ensure that (a) demands are strictly non-negative, and (b) the matrix of compensated price responses is negative semi-definite (property (vi)). We shall call these properties the "general concavity conditions".[5]

Given an estimated preference ordering (or, equivalently, a demand system) the gains and losses resulting from a tax reform are described by the vector of WG values. There are many ways in which information about the reform can be presented. The distribution of gains and losses can be illustrated by a quantile analysis of welfare gains; the change in deadweight loss may be measured by the mean value of welfare gain; and the gains and losses may be aggregated if a social welfare function is specified. For any of these to be meaningful the measure of welfare gain must correspond to a consistent preference ordering for <u>every</u> household in the sample. Outliers cannot be ignored. For example, the change in deadweight loss, which is the mean value of welfare gain, is sensitive to small numbers of households with large absolute values of welfare gain. This is often characteristic of reforms in practice in which most households experience a small gain or loss but in which certain groups, comprising relatively few households, gain or lose substantial amounts. Violations of the general concavity conditions invalidate estimates of deadweight loss. Many empirical studies attempt to overcome this problem by not reporting the mean value of welfare gain but rather the value of welfare gain evaluated at the mean values of prices and incomes in the sample. The general concavity conditions almost always hold at these mean values. Such a procedure may give a misleading picture of the efficiency gains of the reform because welfare gain is a highly nonlinear function of $(\underline{p},y)$. Even for calculations of deadweight loss it is necessary to evaluate WG for the entire distribution of budget constraints observed in the sample.

It follows from this that the conditions for a consistent preference ordering must hold over the relevant domain. In price-income space the relevant domain is the set of all budget constraints implied by the pre- and post-reform vectors of prices and incomes. In quantity space the relevant domain is the set of all consumption bundles that would be chosen by households in the pre- and post-reform allocations. These conditions may be expressed in either primal or dual form.

1. Primal (quantity space)

   The direct utility function $u(\underline{x})$ must be nondecreasing in its arguments (and strictly increasing in at least one) over the relevant domain of $\underline{x}$.

2. Dual (price-income space)

   The general concavity conditions must hold over the relevant domain of $(\underline{p},y)$.

The dual condition was examined above. The primal condition states that nonsatiation holds at each commodity bundle that households consume. This implies that each observed commodity bundle has a positive price support and hence there exists a continuous linear (virtual) budget constraint at which the household would choose to consume the bundle. If the primal condition is not satisfied then no virtual budget constraint (with positive prices) exists. Why does this matter? If welfare is defined in terms of budget constraints then only the general concavity conditions need be satisfied provided that the budget constraints are linear. But if there are any nonlinearities in the budget set, such as a non-negativity constraint on hours worked, then the primal conditions become relevant. The reason is that to evaluate the equivalent income function in the presence of nonlinearities the virtual budget constraint must be computed. If the primal conditions are not satisfied at the value of $\underline{x}$ where the nonlinearity occurs, then with free disposal the assumed preference ordering is inconsistent with $\underline{x}$ having been chosen by the household. Let the set of $\underline{x}$ for which the primal conditions are not satisfied be denoted by X. Then we may say that the preference ordering does not <u>span</u> the set X. Households that are observed to consume at points in X thus create problems for welfare analysis. A satisfactory preference ordering not only satisfies the dual conditions in the relevant price-income domain but also satisfies the primal conditions in the relevant quantity domain.

The primal and dual conditions are <u>not</u> equivalent, and are not implied by each other. It is possible for the general concavity conditions to hold for <u>any</u> strictly positive (p,y) and yet for the primal conditions to fail to hold for some strictly positive $\underline{x}$. To illustrate this point and demonstrate its empirical significance we shall consider some examples from the public finance literature.

4. Example 1 : Labour Supply

Consider the following one-period model of labour supply. Preferences are assumed to be defined over a single composite consumption good and leisure. Labour supply is assumed to be given by the linear function

$$L = b_1 + b_2 w - b_3 y \tag{4.1}$$

where L denotes hours worked, w the real wage rate and y full exogenous income. The price of the consumption good is normalised to unity. Hours of leisure are $(H_M - L)$ where H_M is the total number of hours available for work. The linear labour supply function is chosen here for two reasons. First, it has been used extensively in the analysis of taxation and labour supply (Hausman 1981, Blomquist 1983). Secondly, its simplicity means that it is easy to derive analytically an explicit expression for the virtual budget constraint. The indirect utility function corresponding to (3.9) is[6]

$$v(w,y) = e^{-b_3 w} \left\{ y + \frac{H_M - b_1}{b_3} - \frac{b_2}{b_3} w - \frac{b_2}{(b_3)^2} \right\} \tag{4.2}$$

We now consider the primal and dual conditions for this model. The dual conditions are

(i) f is concave in reference prices. This can be shown to imply that $b_2 - b_3 (H_M - L) \geq 0$ and hence that

$$y \leq \frac{(b_2 - b_3 H_M) + b_1 b_3 + b_2 b_3 w}{(b_3)^2} \tag{4.3}$$

(ii) The demand for leisure is non-negative which implies that

$$y \geqslant \frac{b_1+b_2w-H_M}{b_3} \qquad (4.4)$$

(iii) The demand for the consumption good is non-negative, which implies that

$$y(1-b_3w) \geqslant w(H_M-b_1-b_2w) \qquad (4.5)$$

Figure 2 shows the region in price-income space in which the dual conditions are satisfied. It is bordered by the constraints (4.3)-(4.5). The figure is plotted for US data for 877 male heads of household aged between 29 and 55 taken from the 1980 Michigan Panel Survey of Income Dynamics. Simple linear regressions of hours worked on wages and nonlabour income yield parameter estimates that always violate the general concavity conditions.[7] There are several reasons for this. First, the wage is estimated by the ratio of earnings to reported hours. Measurement error in the latter will lead to biassed estimates of the wage coefficient. Secondly, with nonlinear budget constraints the net wage and income are endogenous. Hence we take Hausman's (1981a) maximum likelihood parameter estimates of the wage and income coefficients (based on a very similar sample of 1085 prime age males in the 1975 PSID) and update them to allow for the increase in wage rates and incomes between 1975 and 1980. The intercept is chosen such that mean predicted and observed hours are equal. This yields the following values, $b_1 = 57.061$, $b_2 = 13.745$, $b_3 = 0.121$.[8] Figure 2 shows also the sample mean wage and full income. They lie in the region within which the dual conditions are satisfied. In fact, for every observation in the sample the observed budget constraint is in the acceptable region.

The primal condition in this static labour supply model is that the virtual budget constraint has a non-negative virtual wage. The virtual wage and virtual income, w* and y* respectively, for a consumption bundle (C,L) are given by the pair of equations

$$L = b_1 + b_2 w^* - b_3 y^* \tag{4.6}$$

$$C = y^* - w^*(H_M - L) \tag{4.7}$$

Hence

$$w^* = \frac{L + b_3 C - b_1}{b_2 - b_3(H_M - L)} \tag{4.8}$$

$$y^* = \frac{b_2 C + (L - b_1)(H_M - L)}{b_2 - b_3(H_M - L)} \tag{4.9}$$

The denominator of these expression is positive if the dual conditions are satisfied. Assuming this to be the case then for the primal condition to hold

$$L \geqslant b_1 - b_3 C \qquad (4.10)$$

Figure 3 shows this constraint in quantity space for the US data used in figure 2. The region that the linear labour supply function cannot span is the triangle bordered by the line drawn in the figure and the two axes. In this region the indifference curves are upward-sloping, and two such are drawn in Figure 3 corresponding to the parameter values assumed above. The direct utility function can be obtained by substituting the expression for w and y given by (4.8) and (4.9) into the indirect utility function (4.2).

In terms of the dual conditions welfare analysis using the linear labour supply function appears straightforward. But this is only part of the story. The primal conditions are violated over a non-negligible region of the commodity space and this restricts the type of problem that can be studied. For example, the model cannot explain non-participation in the labour force (defined as the choice of zero hours of work) for households with small endowments, nor the work decisions of the retired who may choose to work for a few hours a week. Consider a married couple with no nonlabour income and in which the husband's labour supply is fixed at 35 hours per week. The wife's labour supply is variable and the household's exogenous income is the value of the wife's time endowment plus the husband's net of tax earnings. Suppose that the husband's wage was one-half of the mean wage observed in the sample. With the parameters assumed above, zero hours of work would never be chosen by a woman married to such a man if her wage rate was less than 49% of the mean male wage observed in the sample. For such women the virtual wage at zero hours is negative. In certain applications this may be of little consequence. But in general both the primal and dual conditions must be satisfied for the particular reform under consideration.

5. Example 2 : The Demand for Housing Services

The second example concerns the housing market in the UK. The interest in this case is that because of institutional restrictions fostered by government policy observed consumption of housing services is the minimum of demand for housing and the rationed supply which is determined by the policies of public housing authorities. The effective budget constraint is nonlinear and the welfare of a household that is rationed depends upon the virtual rather than the observed prices.

In the UK the price of housing services varies across households both because of differences in tax rates (national and local) and because of the pricing policies of the local housing authorities (King, 1980). Cross-section data can, therefore, be used to estimate both income and price responses. The existence of rationing, however, dampens the effective responses because demands are filtered thorough the rationing mechanism to yield observed consumption levels. For welfare analysis two cases must be distinguished. The first is where the household is unconstrained and money metric utility can be defined over the observed budget constraint. The second is where the household is rationed and money metric utility is defined over the virtual budget constraint computed at the observed consumption bundle. Both primal and dual conditions must be satisfied for welfare analysis to be possible.

To illustrate this we consider a simple model in which preferences are defined over the consumption of housing services and a composite consumption good representing all other commodities. Preferences are described by a flexible functional form and for this example I use my previous estimates of the Deaton-Muellbauer (1980b) almost ideal demand system (AIDS) in which the share of total income devoted to housing is given by

$$s = \beta_1 + \beta_2 \ln y + \beta_3 \ln P_H + \beta_4(\ln P_H)^2 \qquad (5.1)$$

where y denotes real income (deflated by the price of the composite good) and P_H is the real price of housing services. The model was estimated on data for 4,227 households living in England and Wales drawn from the Family Expenditure Survey. The sample period was the fiscal year 1973/4.

The indirecty utility function corresponding to (5.1) is

$$v(P_H, y) = P_H^{-\alpha_1}\{\ln y - \alpha_2 - \alpha_3 \ln P_H - \alpha_4(\ln P_H)^2\} \qquad (5.2)$$

where the parameters of (5.1) and (5.2) are related by

$$\alpha_1 = \beta_2$$

$$\alpha_2 = \frac{-(\beta_3 + 2\beta_4/\beta_2)}{\beta_2^2} - \frac{\beta_1}{\beta_2}$$

$$\alpha_3 = -(\beta_3/\beta_2 + 2\beta_4/(\beta_2)^2)$$

$$\alpha_4 = - \beta_4/\beta_2$$

Consider, first, the dual conditions in the price-income domain. It can be shown that these imply

(i) concavity of f in reference prices:

$$s^2 - s(1-\beta_2) + \beta_3 + 2\beta_4 \ln P_H \leqslant 0 \tag{5.3}$$

(ii) non-negativity of demands:

$$0 \leqslant s \leqslant 1 \tag{5.4}$$

Substitution of (5.1) into the above equations yields the implied relationships between P_H and y. On any realistic scaling these curves, which define the region within which the dual conditions are satisfied, are indistinguishable from the axes. Hence we do not plot the equivalent of Figure 2 for the housing example. The conditions are satisfied for all 4,227 households in the sample. With two commodities the extra flexibility offered by the AIDS functional form has almost no cost in terms of restricting the domain over which the general concavity conditions are satisfied. This illustrates a more general point, namely that the trade-off between flexibility of functional form and consistency of the implied preference ordering is more severe the large the number of commodities.

To investigate the primal conditions we plot the indifference map corresponding to (5.1). There is no explicit form for the direct utility function and its values are computed numerically by solving for the virtual price of housing services and virtual income as functions of quantities and substituting them into the indirect utility function. Although the general concavity conditions are satisfied almost everywhere in the positive orthant of the price-income space, in the quantity space there exists an infeasible region that cannot be spanned by the estimated preference ordering. This is shown in Figure 4. As can be seen from the figure, the indifference curves are spiked-shaped as they approach the line that defines the border of the infeasible region. No point to the northwest of the line in Figure 4 can be supported by the estimated preference ordering. Not only would no household choose to be in this region but there exists no measure of welfare for a household rationed to consume in this region. The indifference map is not defined in the infeasible region. The practical significance of this varies from case to case. Roughly speaking, the infeasible region is where expenditure on housing exceeds one-third of total consumption spending. Of the 4,227 households in the sample only nine were observed to be in the infeasible region. But for the purposes of welfare analysis these households had to be eliminated from the sample, a rather unsatisfactory procedure because of the resulting sample selection bias. In principle, the number of observations that give rise to such problems could be large, and would not be detected by checks on the general concavity conditions alone. Examining the implied indifference map is a useful way of assessing the economic plausibility of the functional form and parameter estimates. This is particularly important when measuring changes in welfare at prices and incomes which are not close to the sample means.

6. Consistent Estimation of Welfare Effects

The main conclusion from the above examples is that to carry out a welfare analysis of a given tax reform we must ensure that the primal and dual conditions are satisfied for <u>all</u> households in the sample. There are two ways in which this might be done. The first is to impose the relevant conditions at the estimation stage. The second is to impose the conditions when carrying out the welfare analysis.

Consider the first method. The difficulty of imposing the relevant conditions depends upon whether the starting point is the specification of a direct or an indirect utility function. Most recent studies have adopted the duality approach in order to simplify the derivation of an observable demand system. One may specify the expenditure function directly or derive it from a demand system chosen such that it may be integrated analytically to yield the indirect utility function using Roy's identity. The great advantage of this approach is that it yields an explicit functional form for the welfare gain. The drawback is that even for quite simple functional forms such as those used in the two examples above, the conditions that must be imposed are complicated functions of both the parameters and the data. This can be seen from equations (4.3)-(4.5) and (5.3) and (5.4). The difficulties of imposing the general concavity conditions globally have been discussed by Christensen and Caves (1980) and Wales (1977) in the context of flexible functional forms. Some advances in methods to impose global concavity have been made by Diewert and Wales (1985) and by McFadden (1985). There remain serious computational difficulties in estimating functional forms that impose the general concavity conditions and yet retain flexibility. Nevertheless as the discussion in McFadden (1985) shows, this is a promising area for research.

Imposition of the primal and dual conditions may be easier if we use the direct utility function. Again this may either be specified directly or obtained by specifying a functional form for the marginal rates of substitution that may be integrated analytically to yield the direct utility function. The latter method was developed by Heckman (1974) for a two-commodity model of labour supply, and was used to assess the impact of child care programmes. To illustrate Heckman's method I shall use a general two-commodity model in which preferences are defined over the consumption levels of two goods, x_1 and x_2. Preferences are described by an underlying direct utility function.

$$u = u(x_1,\ x_2) \qquad (6.1)$$

From the implicit function theorem the marginal rate of substitution between the two commodities, denoted by m, is given by

$$\frac{\partial u/\partial x_2}{\partial u/\partial x_1} = - \frac{dx_1}{dx_2}_{/u=\bar{u}} = m(x_1, x_2) \qquad (6.2)$$

For a suitable specification of the function m the differential equation defined by (6.2) may be solved to give the direct utility function with the constant of integration taken to be the level of utility.

The particular twist that Heckman employs is to label the indifference curves using what is essentially a quantity-metric approach. Choose a fixed value of x_1, x_1^R say, and ask the question, how much of good 2 would a household need in order to be as well off as with the bundle (x_1, x_2)? Denote the answer to this question by x_2^*:

$$u(x_1, x_2) = u(x_1^R,\ x_2^*) \qquad (6.3)$$

Indifference curves are now labelled by the value of x_2^* which is a quantity-metric measure of welfare at the reference quantity x_1^R. In Heckman's empirical application x_1 is hours of leisure, and x_2 is a composite commodity of consumption goods. We may now write the marginal rate of substitution as a function of x_1 and the label x_2^*

$$m = m(x_1, x_2^*) \tag{6.4}$$

This is a method of parameterising the indifference map.

Corresponding to any bundle (x_1, x_2) the value of x^* is given by

$$x_2^* = x_2 + \int_{x_1^R}^{x_1} m(x_1, x_2^*)dx_1 \tag{6.5}$$

The functional form proposed by Heckman is

$$\ln m = \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2^* + \epsilon \tag{6.6}$$

where α_0 may contain observable characteristics that influence preferences and ϵ represents unobservable differences.

$$\int_{x_1^R}^{x_1} m(x_1, x_2^*)dx_1 = \frac{\beta e^{\alpha_2 x_2^*}}{\alpha_1} \left\{ e^{\alpha_1 x_1} - e^{\alpha_1 x_1^R} \right\} \tag{6.7}$$

where $\beta = \exp(\alpha_0 + \epsilon)$

Hence x_2^* is given implicitly by the equation

$$x_2^* = x_2 + \frac{\beta e^{\alpha_2 x_2^*}}{\alpha_1} \left\{ e^{\alpha_1 x_1} - e^{\alpha_1 x_1^R} \right\} \qquad (6.8)$$

In the absence of an explicit solution for x_2^* we cannot integrate (6.6) analytically to obtain an explicit direct utility function. The attraction of the approach, however, is that it is straightforward to see what conditions must be imposed for consistency. These are

(i) $\frac{\partial m}{\partial x_1} \leqslant 0$; for convex preferences

(ii) Equation (6.8) has a unique solution for x_2^*; so that the implied indifference curves do not cross.

(iii) m must be positive for the primal conditions to hold.

An interesting area for further research would be to find ways to generalise Heckman's functional form while satisfying conditions (i), (ii) and (iii).

The demand system is defined implicitly by the equations

$$\beta e^{\alpha_1 x_1} e^{\alpha_2 x_2^*(x_1, x_2)} = \frac{1}{p_1} \qquad (6.9)$$

$$x_2 = y - p_1 x_1 \qquad (6.10)$$

Welfare analysis can be carried out using the quantity-metric measure x_2^* , but the lack of an explicit functional form for either the direct or indirect utility function means that a money metric measure can only be computed by numerical integration.

The sheer computational complexity of imposing the relevant conditions at the estimation stage, especially with heterogeneous preferences, suggests that it may be worthwhile to explore an alternative method. This is to impose the conditions at the stage when welfare analysis is carried out. Where the heterogeneity of preferences is unobservable (a random coefficients model) allowing for such heterogeneity allows us to use extra information in order to impose the conditions implied by theory. In essence the idea is that the primal and dual conditions fail to hold because preferences vary and that the conditions may be imposed by estimating an individual-specific preference parameter vector using the information contained in the estimated residual for the observation. Welfare analysis may then be carried out using the individual-specific measures of welfare gain.

In order to impute an individual-specific effect a loss function must be specified. To illustrate the general principle we shall consider a quadratic loss function in which the objective is to minimise the variance of the prediction of tax revenues. There is no particular reason for this, other than to exploit linear models. A more general approach might be to define the loss function in terms of differences between equivalent income and the expenditure required to purchase the observed consumption bundle at the reference price vector. Moreover, this approach could be extended to the estimation not only of the individual-specific effect but of the mean preference parameter vector as well.

Consider a simple example in which household preferences over commodities may be represented by an expenditure system that is linear in parameters, and in which only linear taxes are considered. Suppose that expenditure on a typical commodity by household h is determined by the model

$$e_h = X_h\beta + \epsilon_h \tag{6.11}$$

This system of expenditure equations may be estimated and the estimates of the parameter vector used to calculate a predicted expenditure after the reform.

$$\hat{e}_h^p = X_h^p \hat{\beta} \tag{6.12}$$

This is the usual procedure when using the parameter estimates to predict the effect of changing the values of the exogenous variables. But although this procedure gives an unbiased estimate of tax revenues it does not give an efficient estimate because it ignores the information contained in the residuals from the fitted regression. To the extent that the residual measures an unobservable household-specific effect it contains information which can be used to reduce the variance of the predictor. A better predictor might therefore be

$$\hat{e}_h^p = X_h^p \hat{\beta} + \hat{\epsilon}_h^o \tag{6.13}$$

where

$$\hat{\epsilon}_h^o = e_h^o - X_h^o \hat{\beta}$$

This estimator (suggested by Feldstein and Taylor (1976) and Feenberg and Rosen (1981) is, however, also inefficient in that it implicitly attributes all of the residual to unobservable household-specific effects. In general, the optimal predictor depends upon the specification of the error term.

The optimal predictor will be taken as that which minimises the variance of the prediction of tax revenues. Total revenues are a linear function of expenditures. If we assume that the error terms in (6.11) are distributed independently across households and across equations (with the exception of that for the n^{th} commodity which is dropped for estimation purposes), then minimizing the variance of the predictor for tax revenues is equivalent to minimizing the variance of the predictor for each expenditure in (6.11). We therefore choose α_h to minimise the variance of the predictor for e_h which we write as

$$\hat{e}_h^p = X_h^p \hat{\beta} + \alpha_h \hat{\epsilon}_h^o \tag{6.14}$$

Note that all three predictors (6.12), (6.13), (6.14) are unbiased, and the differences lie in their relative efficiencies.

We assume that household expenditure may be represented by a mixed error and variance components model. The error term in (6.11) is given by

$$\epsilon_h = X_h\tilde{\beta} + f_h + u_h \qquad (6.15)$$

The first component of the error reflects the fact that preferences vary among the population. The vector $\tilde{\beta}$ is assumed to be distributed with zero mean and covariance matrix Ω. The second component is a household-specific effect which is fixed for each household (hence $f_h^o = f_h^p$) and we assume that these effects are drawn from a distribution with zero mean and variance σ_f^2.[9] The final component measures transitory effects or measurement errors. It is assumed to be identically and independently distributed with zero mean and variance σ_u^2. The three components are assumed to be uncorrelated with each other. With cross-section data it is not possible to estimate σ_f^2, although if we assume that the errors are normally distributed it is possible to obtain maximum likelihood estimates of Ω and the sum of σ_u^2 and σ_f^2. The use of panel data allows the estimation of σ_f^2.

From (6.11), (6.14) and (6.15) the prediction error is

$$e_h^p - \hat{e}_h^p = X_h^p\beta_h + X_h^p\tilde{\beta} + f_h + u_h^p - X_h^p\hat{\beta} - \alpha_h(\epsilon_h^o) \qquad (6.16)$$

The variance of this predictor is

$$V = X_h^pCX_h^{p'} + (X_h^p - \alpha_hX_h^o)\Omega(X_h^p - \alpha_hX_h^o)' + (1-\alpha_h)^2\sigma_f^2 + (1+\alpha_h^2)\sigma_u^2 \qquad (6.17)$$

where C is the covariance matrix of the estimator of the parameter vector β.

The value of α_h which minimizes this variance is given by

$$\alpha_h = \frac{X_h^p \Omega X_h^{o'} + \sigma_f^2}{X_h^o \Omega X_h^{o'} + \sigma_f^2 + \sigma_u^2} \tag{6.18}$$

In other words the optimal predictor of expenditure by household h after the reform is equal to the prediction given by the structural model plus a fraction of the residual for household h from the original regression equal to the proportion of the total variance of the equation attributable to unobservable household-specific effects adjusted for the change in the exogenous variables.

With only cross-section data the value of σ_f^2 cannot be estimated. This leaves two alternatives. Either a value can be imposed using prior information based, perhaps, on other studies which employed panel data (this approach was used by King and Dicks-Mireaux (1980) to estimate permanent income), or the fixed effect can be ignored on the grounds that the most significant household-specific effects are those correlated with the regressors and captured by the specification of random preferences.

From (6.14) the predicted deviation of mean expenditure conditional upon observable characteristics and the value of α_h, may be attributed to individual-specific preferences. This deviation may be spread among the preference parameters according to their relative variances. In practice most estimates of random preference models take Ω to be diagonal. In this way a household-specific preference parameter vector may be obtained, and these parameter estimates used to compute a vector of welfare gains.

Equally, the conditions that are relevant, both primal and dual, are now defined in terms of the household specific preference parameters. To the extent that much of the residual variance in the model with uniform preferences can be accounted for by heterogeneous preferences (including fixed effects) then it is much more likely that the conditions will be satisfied. Indeed, in the example of Section 5, once individual-specific preferences were imputed all of the households in the sample satisfied both primal and dual conditions. But a suggestion for formally imposing the conditions is the following. Once α_h has been determined then spread the predicted residual among the preference parameters subject to the restriction that the necessary conditions hold. Although the implied distribution of the preference parameters will then strictly not be that assumed in estimation, the difference is likely to be small. The example of Section 5 suggests that simply incorporating heterogenous preferences into the welfare measures is likely to be sufficient for all but a very small number of households.

This approach of using the additional information in the observed residual has a good deal in common with the use of shrinkage estimators in statistics, (see Morris (1983) and the accompanying discussion), statistical decision theory, and empirical Bayesian analysis. Because the R^2 in cross-section models is usually low, it is clear that such information is potentially valuable. A full Bayesian treatment poses serious problems becaue it is different to impose the primal and dual conditions on a tractable joint distribution of the parameters of the preference ordering. The prior information contained in the requirement that the consistency conditions be satisfied, is much easier to impose at the stage where we are making inferences about individual-specific parameters. Further use of parametric empirical Bayesian inference appears a promising direction for research.

To illustrate the quantitative importance of individual-specific effects, consider some empirical results for the housing model described in the second example of Section 5. I compared estimation of the welfare gain from abolishing tax relief to home owners both with and without imputed household-specific effects. Allowing for heterogenous preferences reduced the estimated deadweight loss of the tax concessions by between 20 and 25 per cent, but it increased the coefficient of variation of welfare gain by over 40%. Heterogeneity is, therefore, quantitatively important.

7. Conclusions

The stimulus for empirical analysis of tax reforms has come in recent years from proponents of a "supply-side" thesis that there are significant welfare costs to our existing tax system. The issues involved are quantitatively important and the arguments directly impinge on the theoretical and econometric models employed by public finance economists. It is not unreasonable, therefore, to subject the claims made by would-be reformers to serious scrutiny. I interpret this to mean that we need a systematic sensitivity analysis to discover which propositions are robust with respect to changes in the parameterisation (or specification) of the model among which the available data make it difficult to choose. In estimation we do not wish to impose on the data too restrictive a set of assumptions on either the functional form or stochastic distribution. Even if the data are relatively uninformative about certain differences in the parameterisation of the model, it is important to examine the sensitivity of the estimated welfare gains to these differences. The implications of a research programme along these lines are:

(i) To incorporate supply-side behavioural responses the impact of the reform should be measured by the welfare gain, and care be taken to hold constant the reference price vector that used as the basis for comparisons. Much confusion in the literature has been caused by a rather casual approach to the calculation of money measures of gains and losses.

(ii) There is no reason to suppose that there exists a representative consumer, and every reason to suppose that there does not. Hence welfare analysis must be conducted using the full sample of observations.

(iii) Modelling the behaviour of "outliers" is the essence of capturing the effects of tax reform, whereas in conventional econometric estimation outliers are often seen as a potential hindrance to obtaining robust parameter estimates. The top ten percent of the income distribution are often critical to an assessment of the revenue effects of a reform, and the bottom ten percent to the distributional consequences. These households have wage rates, incomes and other characteristics that may be a long way from the mean of the sample.

(iv) Combining (i)-(iii) I have argued that conditional upon the functional form of the model, the implied preference ordering used for welfare analysis must be consistent with the axioms of consumer choice for each observation in the sample. Methods of ensuring this were discussed above. Sensitivity analysis then takes the form of examining the robustness of conclusions about the effects of a reform to alternative functional forms among which it is difficult to discriminate using the available data.[10]

Are violations of either the primal or dual conditions quantitatively significant? In the above examples only simple functional forms were used in order to illustrate the argument and the numbers of violations were small. But with more sophisticated models the number of violations can increase alarmingly. For example, Blundell et al. (1985) in a study of family labour supply in the UK that investigated the role of demographic characteristics in great detail found large numbers of violations of the general concavity conditions for even the most flexible functional form estimated. These occurred for between 16% and 26% of the sample. As Blundell and Meghir (1985) themselves comment "there is clear evidence of an underlying trade-off between flexibility and theory consistency." Similarly, in studies of consumer demand using large cross-section data sets Baccouche and Laisney (1985) and Hughes (1985) both found large numbers of violations. Hughes (1985) discovered that different flexible functional forms could be ranked in terms of the number of violations. Performance in this dimension might be used as an informal criterion of model selection.

Does it really matter if the primal or dual conditions are violated by some households in the sample? Surely, the models that we estimate are no more than approximations (or to use Leamer's (1985) phrase "economic metaphors") to the underlying model of economic behaviour, and it should come as no surprise when some observations appear to be inconsistent with the theory. This misses the point. The fact that our models capture only some of the many influences of policy variables on behaviour is certainly important and implies the need for sensitivity analysis. But for any given parameterisation of the model the analysis must be consistent with the concept that is being measured. Violations of either the primal or dual conditions mean that the estimated welfare gains for those households are meaningless and can lead to severe bias in the estimate of statistics such as the mean welfare gain (that is, the change in deadweight loss).[11] The conventional approach to this problem is simply to drop these households from the sample when doing welfare analysis. It is clear that this is not an attractive procedure. There is no obvious way to correct for the resulting sample selection bias because the model provides no means of relating the welfare gain of the excluded households to the estimated gain of the included households. The selection criterion is the result of the failure of the model. Moreover, the number of households dropped because of violations will vary with the choice of functional form thus making it more difficult to carry out sensitivity analysis on a uniform sample.

Perhaps the following practical argument will convince those readers who are as yet unpersuaded. Household expenditure surveys are used by governments to predict who will gain and who will lose from tax changes. It is unlikely that the Chancellor of the Exchequer could stand up in the House of Commons and announce that according to Treasury calculations almost every family would benefit from his proposals, except that is for the 20 per cent of families who unfortunately were deleted from the analysis because their behaviour appeared to be inconsistent with the axioms of consumer choice. A promising direction for future research is the imposition of the primal and dual conditions by allowing for variations in preferences among the population.

In the last resort of course many will argue that it is fruitless to analyse policy changes, that economists are either ignored or used by policy-makers. But the fact that empirical analysis may have some impact on decisions means that, even if our present knowledge is very limited, we should strive to report that knowledge in a manner that is credible, stressing those results that seem to be most robust.

FOOTNOTES

1. Similarly, in an intertemporal setting equivalent wealth may be defined as that value of initial full wealth which, at the reference price vector (including interest rates), affords the same level of expected utility as can be attained given the actual budget constraint and distribution of future wages and prices.

2. In an economy of identical households, the optimum tax rates are obtained by maximising (3.5) subject to the government's revenue constraint. It can be shown (King, 1983a) that this is equivalent to minimising a measure of deadweight loss. When households are heterogeneous the optimum depends upon explicit interpersonal comparisons.

3. The term "reform" is sometimes used in the literature to denote only a local change from some given initial position. The analysis of such marginal reforms requires much less information than in the case of discrete changes, as Ahmad and Stern (1983) have shown. Knowledge of aggregate responses alone provides sufficient information to evaluate the reform. In this paper we focus on the analysis of non-marginal reforms.

4. The problem is even more severe in the case of the compensating variation which would, for example, incorrectly compare a set of mutually exclusive reforms at reference prices corresponding to each of the alternatives. (Kay 1980, King 1983a). For this application the equivalent variation would be appropriate because it would employ a common reference price vector given by the pre-reform prices.

5. These are called "regularity conditions" by Christensen and Caves (1980).

6. See Hausman (1981a). We define y to be full exogenous income in the static model of labour supply, whereas Hausman takes it to be nonlabour income. This is of no consequence for our analysis. In an intertemporal model y represents full expenditure in the period, and with preferences that exhibited intertemporal separability (4.2) would be the conditional indirect utility function.

7. In my experience this is true also for similar samples drawn from the UK Family Expenditure Survey.

8. Hausman's (1981a) wage coefficient was 0.0113, and for the income coefficient we take the mean value of the estimated truncated normal distribution assumed for this coefficient of -0.153. These were converted to 1980 weekly values and adjusted for the fact that we use full income rather than virtual nonlabour income.

9 . We take a random effects rather than a fixed effects model because the first component of the error term already allows for a household-specific effect that is correlated with the regressors.

10. The computational burden of a sensitivity analysis of this kind may be greatly lessened by the use of a package such as TRAP (King 1983a).

11. If the primal conditions are violated then no welfare measure for the household may exist (as in the infeasible region in Figure 4). When the dual conditions are violated then the compensated own-price responses may be positive, and the introduction of a distortionary tax may raise predicted welfare and lower estimated deadweight loss.

REFERENCES

Ahmad, E. and Stern, N.H. (1984), "The Theory of Reform and Indian Indirect Taxes", Journal of Public Economics 25, 259-298.

Ashenfelter, O. and Heckman, J.J. (1974), "The Estimation of Income and Substitution Effects in a Model of Family Labor Supply", Econometrica 42, 73-85.

Ashworth, J.S. and Ulph, D.T. (1981), "Estimating Labour Supply with Piecewise Linear Budgets, in (eds.) Brown, C.V. Taxation and Labour Supply Allen and Unwin, London.

Atkinson, A.B. (1985), "How Should We Measure Poverty? : Some Conceptual Issues,", mimeo, LSE.

Auerbach, A.J. and King, M.A. (1983), "Taxation, Portfolio Choice and Debt-Equity Ratios: A General Equilibrium Model", Quarterly Journal of Economics 98, 587-609.

Baccouche, R. and F. Laisney (1985), "Analyse Microeconomique de la Reforme de la TVA de juillet 1982 en France : Premiers Resultats". mimeo, Toulouse.

Blomquist, N.S. (1983), "The Effect of Income Taxation on the Labour Supply of Married Men in Sweden", Journal of Public Economics 22, 169-197

Blundell, R. and Meghir, C. (1985), "Selection Criteria for a Microeconometric Model of Labour Supply", Centre for Economic Policy Research Discussion Paper No. 57, mimeo, London.

Blundell, R. Meghir, C., Symons, E. and Walker I. (1985), "Alternative Specifications of Labour Supply and the Simulation of Tax Reforms", mimeo, University College London.

Burtless, G. and Hausman, J.A. (1978), "The Effect of Taxes on Labour Supply: Evaluating the Gary Negative Income Tax Experiment", Journal of Political Economy 86, 1103-1130.

Christensen, L. R. and D. W. Caves (1980), "Global Properties of Flexible Functional Forms", American Economic Review 70, 422-432.

Deaton, A.S. (1981), "Optimal Taxes and the Structure of Preferences", Econometrica 49, 1245-1260

Deaton, A.S. (1981), "Econometric Issues for Tax Design in Developing Countries", in (eds.) Newbery, D. and Stern, N.H.

Deaton, A.S. and J. Muellbauer, 1980a, Economics and Consumer Behaviour (Cambridge University Press)

Deaton, A.S. and J. Muellbauer, 1980b, An Almost Ideal Demand System, American Economic Review 70, 312-326.

Debreu, G. 1951, The coefficient of Resource Utilization, Econometrica 19, 273-292.

Debreu, G. 1954 A Classical Tax-Subsidy Problem, Econometrica 22, 14-22.

Diewert, W. E. (1981), "The Measurement of Deadweight Loss Revisited", Econometrica 49, 1225-1244.

Diewert, W. E. and T. J. Wales (1985), "Flexible Functional Forms and Global Curvature Conditions", University of British Columbia Discussion Paper No. 85-19, mimeo.

Feenberg, D. R. and H. S. Rosen (1983), "Alternative Tax Treatments of the Family : Simulation Methodology and Results", in (ed.) Feldstein, M.S. Behavioural Simulation Methods in Tax Policy Analysis, Chicago University Press.

Feldstein, M. S. and A. K. Taylor (1976), "The Income Tax and Charitable Contributions", Econometrica 44, 1201-1222

Gorman, W.M., 1961, On a Class of Preference Fields, Metroeconomica 13, 53-56.

Hausman, J.A. (1981a) "Labor Supply" in (eds.) Aaron, H.J. and Pechman, J.A. How Taxes Affect Economic Behavior, Brookings Institution.

Hausman, J.A. (1981b), "Exact Consumers' Surplus and Deadweight Loss", American Economic Review 71,662-676.

Hausman, J.A. (1983a), "Stochastic Problems in the Simulation of Labor Supply", in (ed.) Feldstein, M.S. Behavioral Simulation Methods in Tax Policy Analysis, Chicago University Press.

Hausman, J.A. (1983b), "Taxes and Labor Supply", National Bureau of Economic Research Working Paper No. 1102.

Hausman, J.A. (1985), "The Econometrics of Nonlinear Budget Sets", Econometrica, 53, 1255-1282.

Hausman, J.A. and W. E. Taylor (1981), "Panel Data and Unobservable Individual Effects", Econometrica 49, 1377-1398.

Heckman, J.J. (1974), "Shadow Prices, Market Wages, and Labor Supply", Econometrica 42, 679-694.

Heckman, J.J. (1983), "Comment on Hausman", in (ed.) Feldstein, M.S. Behavioral Simulation Methods in Tax Policy Analysis, Chicago University Press.

Heckman, J.J. (1974), "Effects of Child Care Programs on Womens' Work Effort", Journal of Political Economy. 82, S136-163.

Hughes, G. A. (1985), "Concavity and the Estimation of Consumer Demand Equations", mimeo, Edinburgh University.

Hurwicz, L. and Uzawa, H. (1971), On the Integrability of Demand Functions, in: J. Chipman, L. Hurwicz, M. Richter and H. Sonnenschein, eds., Preferences, Utility and Demand, (Harcourt Brace Jovanovich, New York).

Kay, J.A., 1980, The Deadweight Loss from a Tax System, Journal of Public Economics 13, 111-119.

King, M.A., 1980, An Econometric Model of Tenure Choice and Demand for Housing as a Joint Decision, Journal of Public Economics 14, 137-159.

King, M.A. (1983a), "Welfare Analysis of Tax Reforms Using Household Data", Journal of Public Economics 21, 183-214.

King, M.A. (1983b), The Distribution of Gains and Losses from Changes in the Tax Treatment of Housing, in: M.S. Feldstein, ed., Behavioural Simulation in Tax Policy Analysis (Chicago University Press).

King, M.A. (1983c), "Tax Policy and Family Welfare", Second Departmental Lecture on Economic Policy, University of Southampton, mimeo.

King, M. A. and L-D.L. Dicks-Mireaux (1982), "Asset Holdings and the Life-Cycle", Economic Journal 92, 247-267.

Leamer, E.E. (1985), "Sensitivity Analyses Would Help", American Economic Review 75, 308-313.

McFadden, D. (1985), "Specification of Econometric Models", Presidential Address to the Fifth World Congress of the Econometric Society.

McKenzie, L., (1956), Demand Without a Utility Index, Review of Economic Studies 24, 185-189.

Miller, M.H. (1977), "Debt and Taxes", Journal of Finance 32, 261-275

Moffitt, R. (1985), "The Econometrics of Piecewise-Linear Budget Constraints : A Survey and Exposition of the Maximum Likelihood Method", Brown University Working Paper No. 85-10, mimeo.

Morris, C. N. (1983), "Parametric Empirical Bayes Inference : Theory and Applications", Journal of the American Statistical Association, 78 47-55.

Muellbauer, J. (1975), Aggregation, Income Distribution and Consumer Demand, Review of Economic Studies 42, 525-543.

Muellbauer, J. (1976), Community Preferences and the Representative Consumer, Econometrica 44, 972-999.

Musgrave, R.A., Case, K.E. and Leonard, H. (1974), "The Distribution of Fiscal Burdens and Benefits", Public Finance Quarterly 2, 259-311.

Pechman, J.A. (1985), Who Paid the Taxes, 1966-85? Brookings Institution, Washington, DC.

Pechman, J.A. and Okner, B.A. (1974), Who Bears the Tax Burden?, Brookings Institution, Washington DC.

Rothbarth, E. 1941, The Measurement of Changes in Real Income Under Conditions of Rationing, Review of Economic Studies 8, 100-107.

Samuelson, P.A. 1974, Complementarity - An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory, Journal of Economic Literature 12, 1255-1289.

Stern, N.H. (1984), "On The Specification of Labour Supply Functions", ESRC Programme on Taxation, Incentives and the Distribution of Income Discussion Paper No. 50, mimeo LSE.

U.S. (1985), The President's Tax Proposals to the Congress for Fairness, Growth and Simplicity, US Government Printing Office, Washington DC.

Varian, H. (1984), Microeconomic Analysis (2nd edition), W.W. Norton and Co., New York.

Wales, T.J. (1977), "On the Flexibility of Flexible Functional Forms : An Empirical Approach", Journal of Econometrics 5, 183-193.

Wales, T.J. and Woodland, A.D. (1980), "Sample Selectivity and the Estimation of Labour Supply Functions", International Economic Review 21, 437-468.

Weisbach, M.S. (1985), "The Economic Effects of the Earned Income Tax Credit", mimeo, MIT.

Zabalza, A. (1983), "Optimization Errors and Measures of Welfare Change", Centre for Labour Economics Discussion Paper No. 176, mimeo LSE.

Zabalza, A. and Arrufat, J.L. (1983), "Efficiency and Equity Effects of Reforming the British System of Direct Taxation: A Utility Based Simulation Methodology", Centre for Labour Economics Discussion Paper No. 175, mimeo LSE.

Figure 1:Budget Constraint,UK1985

Married man;constant MPL

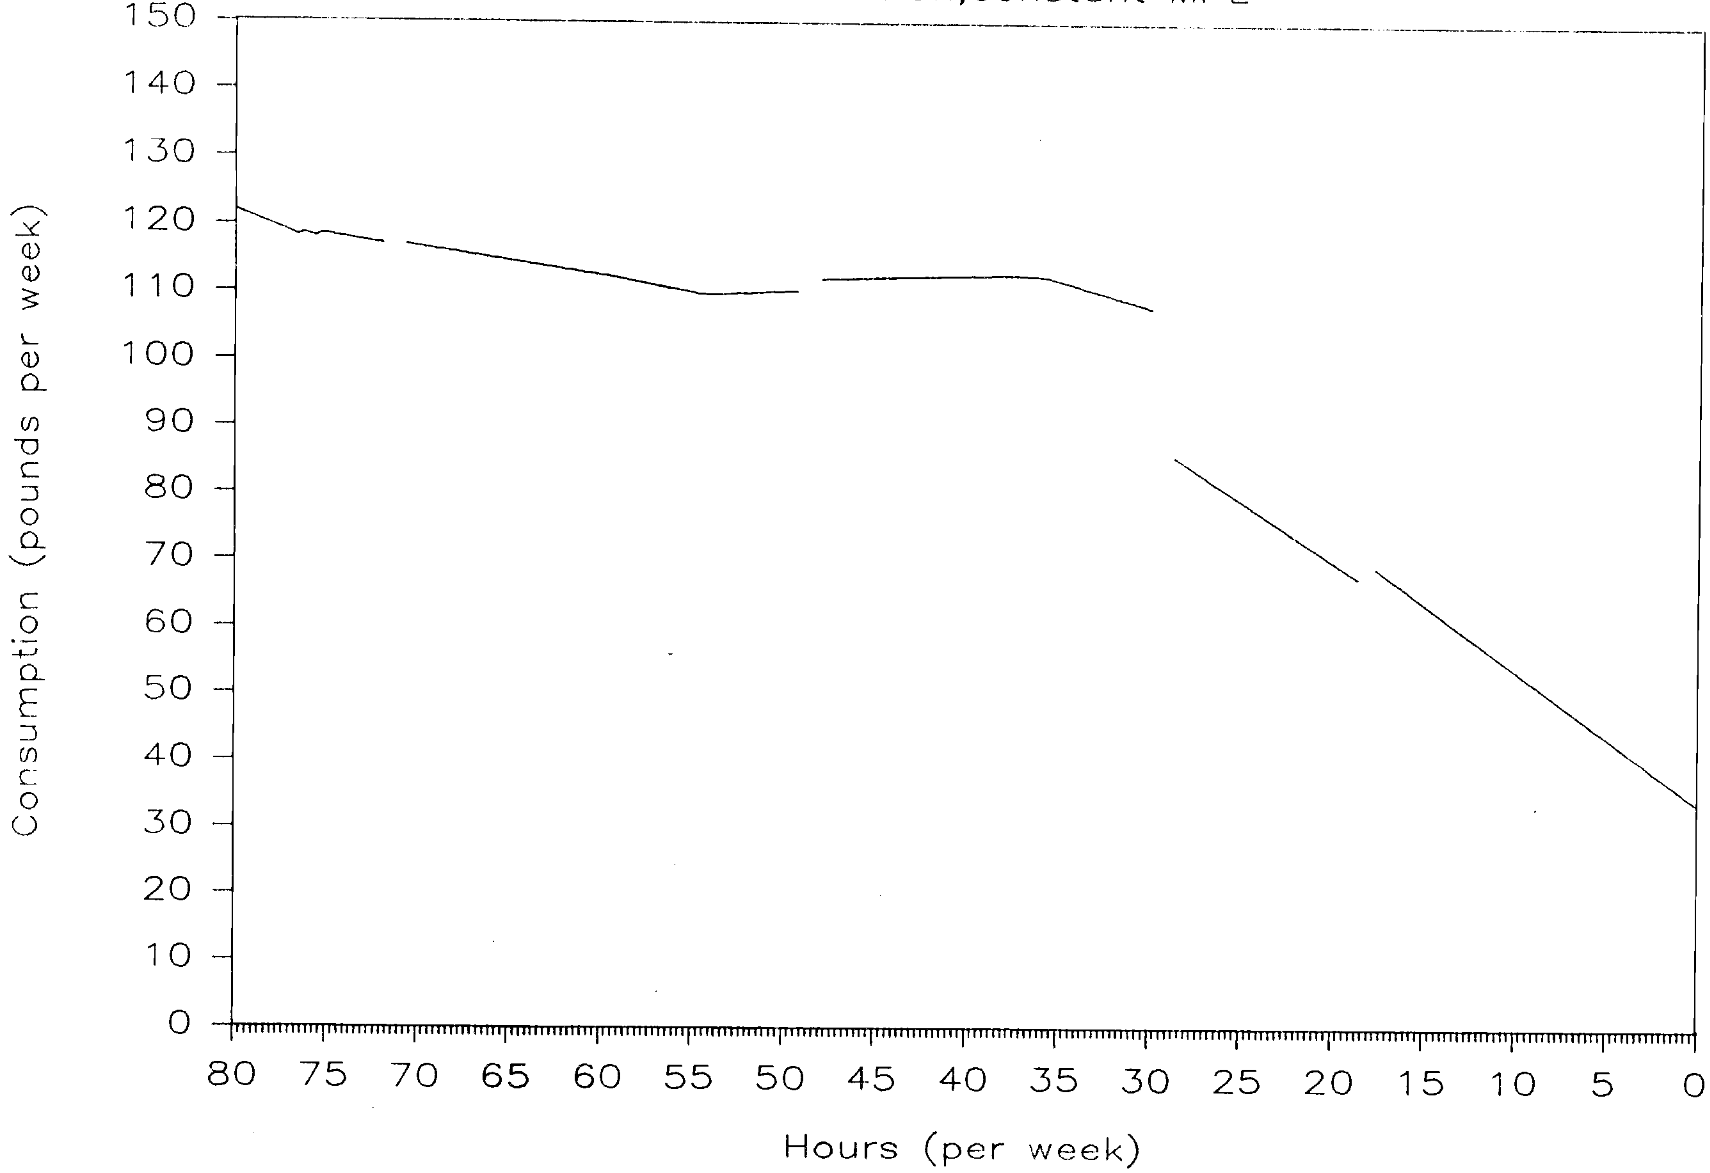

Figure 2: Dual Conditions

Linear Labour Supply Function

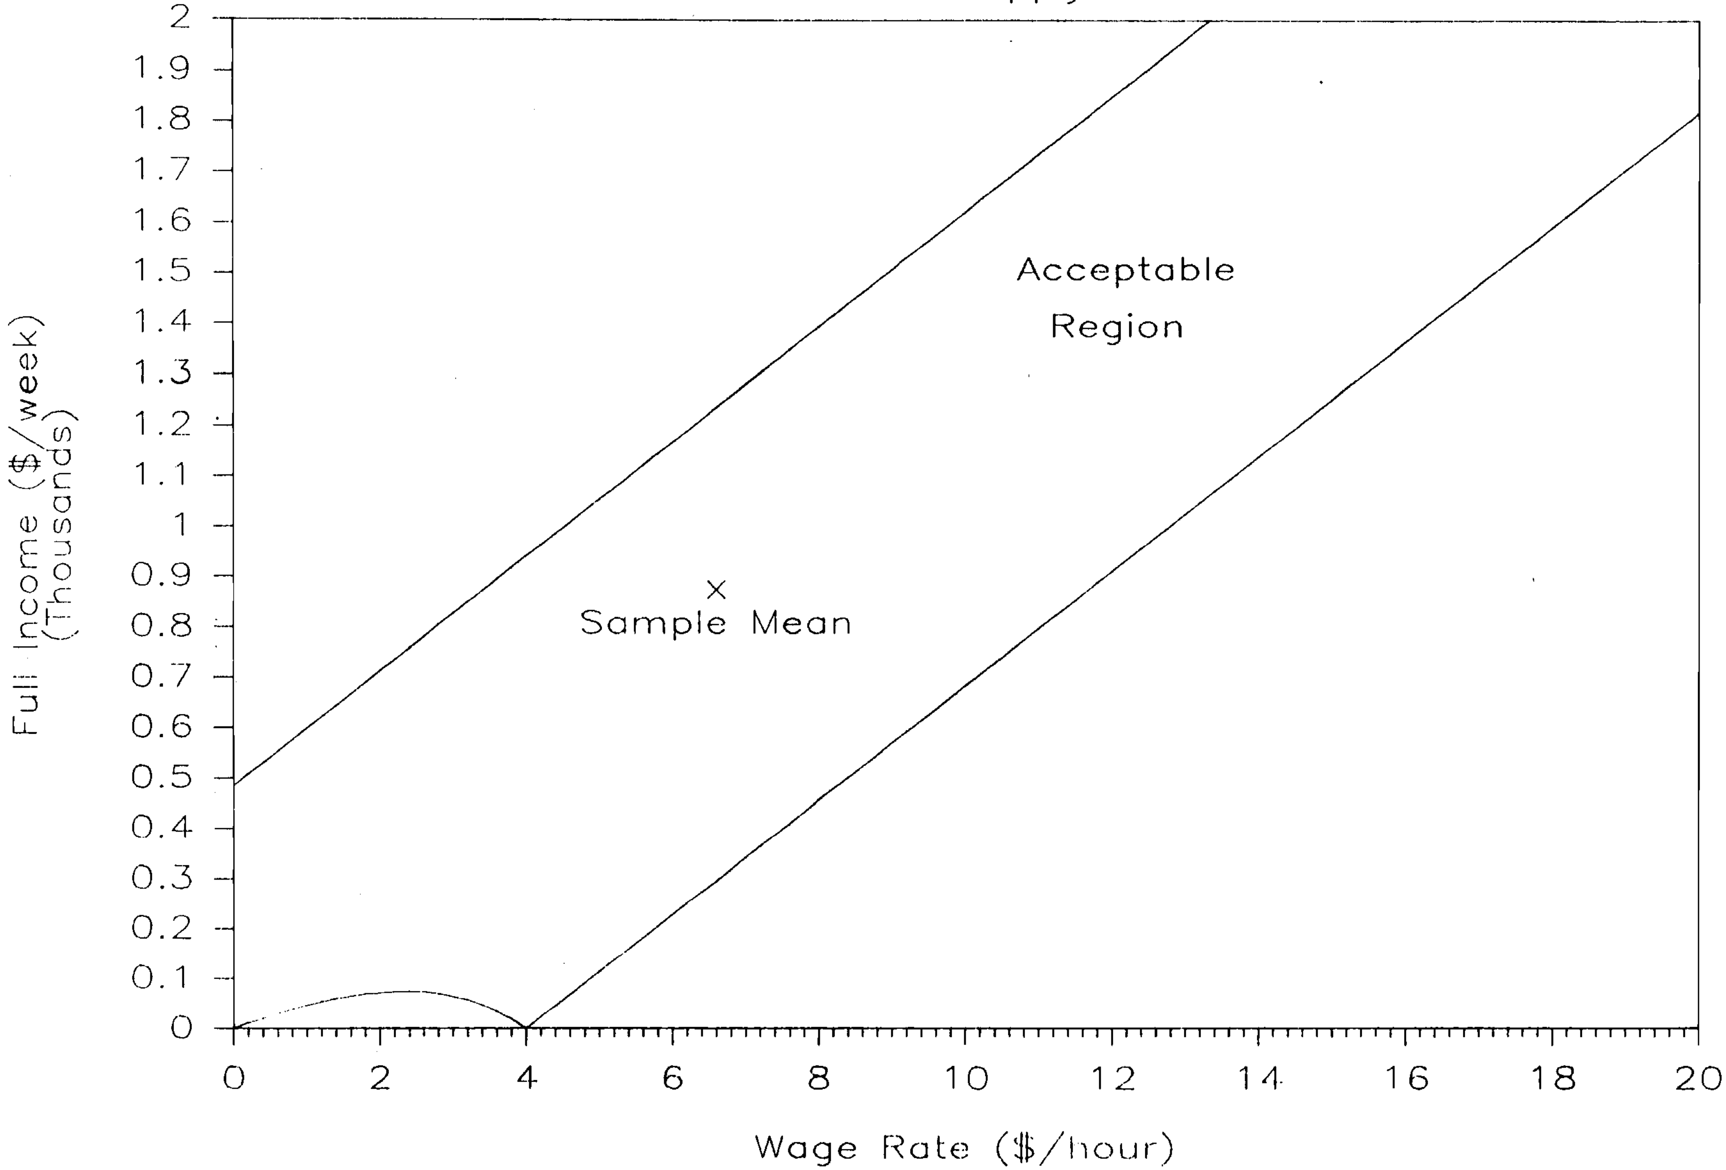

Figure 3: Primal Conditions

Linear Labour Supply Function

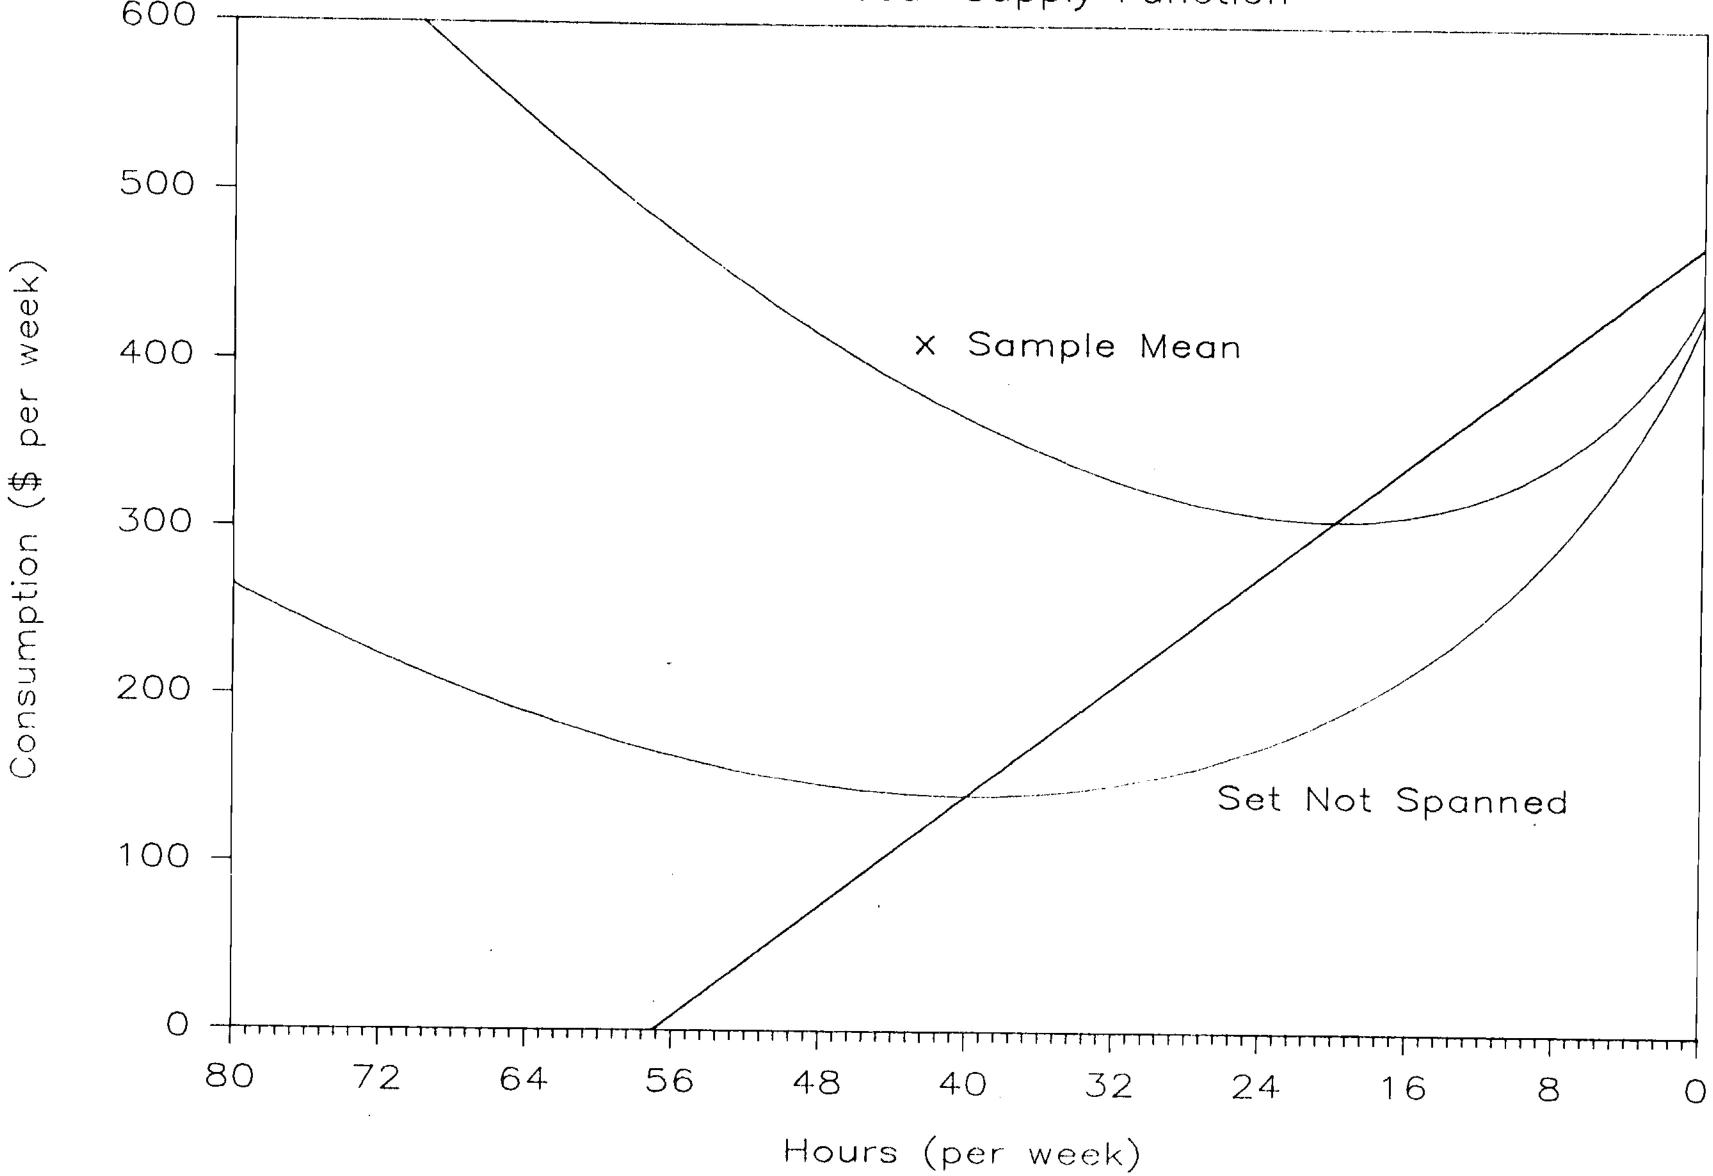

Figure 4: Primal Conditions

AIDS Housing Demand Function