Journal of Public Economics 36 (1988) 23–52. North-Holland

LABOUR SUPPLY SPECIFICATION AND THE EVALUATION OF TAX REFORMS

Richard BLUNDELL and Costas MEGHIR
University College London, Gower Street, London WC1E 6BT, UK, and Institute for Fiscal Studies

Elizabeth SYMONS
University College of Wales, Aberystwyth, Dyfed SY23 3DB, UK, and Institute for Fiscal Studies

Ian WALKER*
University of Manchester, Manchester M13 9PL, UK, and Institute for Fiscal Studies

Received July 1985, revised version received January 1988

This paper investigates the practical importance of the functional specification of labour supply equations for the analysis of tax/benefit reform. We consider two labour supply specifications, one derived from Stone–Geary preferences and the other from a generalisation that relaxes some of the more critical restrictions in that model. We use these equations to contrast the effects of two reforms to the U.K. income tax system; one involves income effects almost exclusively while the other implies a significant number of large increases in marginal tax rates.

1. Introduction

Considerable empirical research has been conducted into the determinants of individual labour supply decisions. A major motivation for this research has been concern over the effects of income taxation and income support schemes on labour market behaviour. However, as many of these studies simply report estimated summary elasticity measures of labour supply behaviour, it is not usually possible to infer the consequences of the estimates for tax policy reform. In turn this makes it difficult to compare the policy implications of different labour supply specifications. There are two main reasons for this. First, taxes and benefits usually result in budget constraints that are nonlinear and often nonconvex, so that marginal tax rates differ markedly across individuals. As Hausman (1985) has emphasised, in such circumstances a change in either the gross wage rate or some parameter of

*We are grateful to the Institute for Fiscal Studies for allowing us to incorporate their tax and benefit program within the simulation routine used in this paper, and for providing the data. We would like to thank Nicholas Stern, two anonymous referees and seminar participants at Bristol, Oxford and the London School of Economics for their comments and advice. All remaining errors are ours alone. The finance for this research was provided by ESRC under grants B2002 2060 and B0225 0004.

0047–2727/88/$3.50

the tax system may cause individuals to shift from one segment of the budget constraint to another. These movements cannot be captured without a detailed knowledge of the labour supply curve for each individual in any sample to be used for tax reform analysis. Secondly, even where sufficient detail is provided, if empirical results are not consistent with economic theory across the whole sample, utility comparisons between distinct points in the budget set, which may be necessary for tax policy simulation, cannot be made.

The analysis of the welfare and behavioural effects of income taxes and benefits therefore requires the choice of some specific representation for preferences either explicitly in the guise of a direct (or indirect) utility function or implicitly as a labour supply function. This choice of specification is usually dictated by econometric convenience subject to some minimum degree of 'flexibility'. Since attention has traditionally focused on the econometric difficulties associated with 'corner' solutions arising from either nonparticipation or from nonlinearities in the post-tax budget constraint, studies of tax reform have normally employed labour supply equations that either correspond to some known form of utility function or can be easily inverted so as to obtain the (virtual) wage that supports any given level of labour supply. In the former category the Stone–Geary utility function and the constant elasticity of substitution (CES) utility functions have proved popular. The Stone–Geary function yields an earnings equation that is linear in the net marginal wage and intercept income variables [see Ashenfelter and Ham (1979) and Hurd and Pencavel (1981)], while the CES results in the log of the ratio of leisure to consumption being linear in the log wage [see Arrufat and Zabalza (1986)]. In the latter category linear or log-linear forms have proved popular, as in Burtless and Hausman (1978) and Hausman (1985), for example. Their properties in simulation are discussed in King (1987).

Where a single labour supply decision is under analysis, the correspondence between the labour supply function and the underlying form for preferences is easily determined. Indeed, Stern (1986) provides a detailed exposition of the underlying properties and theoretical restrictions of a number of popular specifications. As the focus of the present paper is to address the use of labour supply models for policy analysis we choose our labour supply specifications so as to highlight the following three problems: the use of estimates to make predictions of labour supply responses to taxes, the difficulty of achieving a model of behaviour that is consistent with the economic framework, and the implications of the restrictiveness of the chosen functional form.

In Blundell and Meghir (1986), a variety of labour supply specifications were fitted to the observed behaviour of a sample of married women in the U.K. 1981 Family Expenditure Survey. These specifications generalised the simple Stone–Geary linear earnings equation by allowing for general interactions between income, wage, and demographic variables. Using conventional econometric tests the more restrictive forms that imply additive preferences and a labour supply equation that is monotonic in the wage were rejected by the data. However, the extent to which such statistically unacceptable restrictions produce misleading policy predictions depends on the properties of the models under simulation. In order to illustrate the degree of sensitivity of policy conclusions to the specification of labour supply we use two alternative female labour supply equations and a large sample of households to simulate the effects of inequality, labour supply decisions, and efficiency of two reforms to the U.K. tax and benefit system.

In section 2 of the paper we outline the properties of the labour supply equations and explain the methodology used in simulating labour supply responses to tax and benefit reforms. We are inevitably led to the problem posed by estimated models which, for certain individuals, violate the conditions imposed by economic theory. We consider the effects of tax reform only on the labour supply of married women because our research on household labour supply [see, for example, Blundell and Walker (1986) and Meghir (1985)] indicates that the male hours of work is not a dimension of household behaviour that exhibits significant sensitivity of economic variables. This finding is not unusual and indeed is a main conclusion of the exhaustive survey of the literature in Pencavel (1987).

In section 3 we explain the nature of the two tax reforms. These two reforms are useful examples since one involves income effects almost exclusively while the other involves a substantial number of large increases in marginal tax rates. The reforms relate to the taxation of the incomes of married couples versus individuals which has been an area of significant controversy in the United Kingdom, the United States, and elsewhere. We provide the results from applying our simulation methodology to the alternative labour supply specifications and highlight their differences for each of the reforms. We concentrate on the consequences of reform for government revenue, the distribution of welfare across households, and the predicted labour supply responses.

2. Simulating labour supply effects

Although taxes and benefits generally cause budget constraints to be nonlinear, Hausman (1979) has shown that where the resulting budget constraint is convex and where the integrability conditions are satisfied, predictions from the labour supply equation alone are sufficient to evaluate the behavioural effects of policy-induced changes in the constraints. In practice, however, means tested benefits such as Housing Benefit (HB) and Family Income Supplement (FIS) in the United Kingdom, or Aid for

Families with Dependent Children (AFDC) in the United States lead to significant nonconvexities in budget constraints for some individuals. Moreover, in the United Kingdom, the National Insurance Contribution (NIC) system (the equivalent to social security contributions in the United States) also leads to nonconvexities and discontinuities that affect the budget constraints of all individuals.

With nonconvexities and discontinuities in the budget constraint, welfare comparisons across distinct points on the budget constraint are required so as to determine the global utility-maximising solution. If the explicit form of the direct utility function is not known, the computation of utility across different points on the budget constraint involves using the appropriate support wage rates and income levels to calculate indirect utility in order to make these welfare comparisons. If the indirect utility function is known explicitly, this simply involves evaluating it at the appropriate wage rate and income. Thus, when the utility at a kink in the budget constraint is required, the relevant wage is that which would support that kink as an interior maximum. In certain cases, for example the linear labour supply curve or the linear earnings equation, there will be an analytical solution for this virtual wage supporting any point on the budget constraint. In other cases numerical methods will be needed.

The more restrictive of the two specifications we adopt in this paper is, indeed, the linear earnings equation derived from the Stone–Geary utility function, while the more complicated case has a labour supply curve which is not necessarily monotonic in the wage. These are outlined in subsection 2.1 below. In subsection 2.2 we explain the detail of the simulation methodology and in subsection 2.3 we highlight the issues that need to be faced when using estimated labour supply models to conduct policy work of this nature.

2.1. The properties of the labour supply models

The labour supply equations we use are generated from the class of preferences over household consumption and the allocation of the wife's time given by the indirect utility function:

$$V=\frac{\mu+a(w,p,z)}{b(w,p,z)}, \tag{1}$$

where w is the marginal wage rate, p is the price level, z is a vector of household characteristics, and μ is intercept income (i.e. μ is the level of net household income corresponding to the intercept of the linearised budget constraint). The forms of $a(\cdot)$ and $b(\cdot)$ are given by:

$$a(w,p,z)=\alpha_f(z)w-2\alpha_{fq}(z)w^{1/2}p^{1/2}-\alpha_q(z)p, \tag{2}$$

$$\ln b(w,p,z)=(\beta_f(z)-e)\ln w+(\beta_q(z)+e)\ln p+1/2\beta_{ff}(\ln(w/p))^2, \tag{3}$$

where the subscript f denotes female time and the subscript q denotes household goods consumption, e is the stochastic component of preferences assumed to be $N(0,\sigma^2)$, and $\beta_f+\beta_q=1$ ensures the zero degree homogeneity of the indirect utility function. In the flexible functional form literature [see Diewert (1974)], (2) corresponds to a Generalized Leontief specification, while that of (3) corresponds to a Translog.

For the Stone–Geary specification corresponding to (1)–(3) both the α_{fq} and β_{ff} terms in (2) and (3) would be omitted. In this case $a(\cdot)$ would be of the Leontief form so that $\alpha_f(z)$ would represent the *maximum* time available for work while $\alpha_q(z)$ would represent the *minimum* level of consumption. In general $\mu+a(\cdot)$, sometimes termed supernumerary income,[1] is best interpreted as the total value of resources which the household can allocate at its discretion. Although some properties of the Stone–Geary model are restrictive and will be relaxed in our more general specification, it seems useful to choose a form for indirect utility that retains this underlying feature of the Stone–Geary model. In (1) the function $b(\cdot)$ can be interpreted as the price index which deflates $\mu+a(\cdot)$, a money measure of resources, into utility terms. Indeed, in the Stone–Geary case $b(\cdot)$ is a simple Cobb–Douglas index and $\beta_f(z)-e$ would be interpreted as the marginal value of female time. Just as in our discussion of $\alpha_f(z)$, we may expect this marginal value to depend on household characteristics z as well as varying randomly across the population with e. Substituting (2) and (3) into (1) and applying Roy's Identity yields a labour supply equation which is most conveniently written as the earnings equation:

$$wh=\alpha_f(z)w-\alpha_{fq}(z)w^{1/2}-(\beta_f(z)+\beta_{ff}\ln w-e)(\mu+a(\cdot)), \tag{4}$$

where w and μ are now defined in real terms.

In the empirical specification of (4), the dependence of $a(\cdot)$ and $b(\cdot)$ on the vector of household characteristics z can be explained in the following way. We allow both the α's and β_f to depend on the number of children n_i, $i=1,2,3$, in each of the three age groups 0–4, 5–10, 11+. In addition, we allow β_f to depend on the age (A) and education (E) of the wife. To capture the dominant effect of the youngest child on female labour market decisions, we also define three zero–one dummy variables D_1, D_2, and D_3 which

[1] Note that $a(\cdot)$ subsumes the parameter representing the maximum *leisure* time available. Indeed, defining this maximum time to be T, the term $\mu+a(\cdot)$ in (1) can be rewritten in more standard notation as $(\mu+Tw)-[(T-\alpha_f)w-2\alpha_{fq}w^{1/2}p^{1/2}-\alpha_q p]$. The first term in square brackets is often referred to as full income.

indicate the age group containing the youngest child. The resulting specifications take the form:

$$\alpha_f(z)=\alpha_f^0-\alpha_f' n_1-\alpha_f'' D_2-\alpha_f''' D_3, \tag{5a}$$

$$\alpha_{fq}(z)=\alpha_{fq}^0+\alpha_{fq}' n_1+\alpha_{fq}''(D_2 n_2)^{1/2}+\alpha_{fq}'''(D_3 n_3)^{1/2}, \tag{5b}$$

$$\alpha_q(z)=\alpha_q^0+\alpha_q' n_1+\alpha_q'' n_2+\alpha_q''' n_3, \tag{5c}$$

$$\beta_f(z)=\beta_f^0+\beta_f' D_1+\beta_f'' D_2+\beta_f''' D_3+\beta_f^a A+\beta_f^{aa} A^2+\beta_f^e E. \tag{5d}$$

The model description is completed by choosing a distribution for e. It should be noted that although we choose a normal distribution this assumption on e as well as the linearity of (4) in μ were rigorously tested in the Blundell and Meghir (1986) study.

As described above, the general form given by (4) breaks additive separability of preferences in two ways: first via α_{fq} in $a(\cdot)$ and secondly via β_{ff} in $b(\cdot)$. If both of these parameters were restricted to zero the labour supply equation would become that associated with the Stone–Geary utility function, i.e. the linear earnings system. Blundell and Meghir (1986) estimate the general form given by (4) and a variety of nested cases. In all estimated models supernumerary income, $\mu+a(\cdot)$, is positive for all sample points. The LES specification, with $\alpha_{fq}(z)=\beta_{ff}=0$, was rejected on a likelihood criterion against the general model. The restriction $\alpha_{fq}=0$ alone could not, however, be rejected. Moreover, the tests on the normality assumptions were decisively rejected for the LES model. Thus, here we investigate the model with $\alpha_{fq}=0$ and β_{ff} nonzero, which we refer to as the GPF model since it stems from the Gorman Polar Form of preferences in (1). This is compared to the model with $\beta_{ff}=\alpha_{fq}(z)=0$ which we refer to as the LES case since it implies that earnings are a linear function and which is derived from Stone–Geary preferences. The estimated parameters of both models are presented in appendix A.

In terms of their economic properties, perhaps the most important distinction between the LES and GPF labour supply equations is that the LES restricts labour supply to be a monotonic function of wage, while the GPF allows labour supply to take various forms. For example, we may expect labour supply to be forward sloping at low wages and backward bending at higher wages. The effect of the parameter restrictions on the slope of the labour supply curve can be seen from the wage derivative. From (4) we have:

$$\partial h/\partial w=\beta(\mu-\alpha_q)w^{-2}+(\tfrac{1}{2})\alpha_{fq}(1-2\beta)w^{-3/2}-\beta_{ff}(\mu+\mu(\cdot))w^{-2}, \tag{6}$$

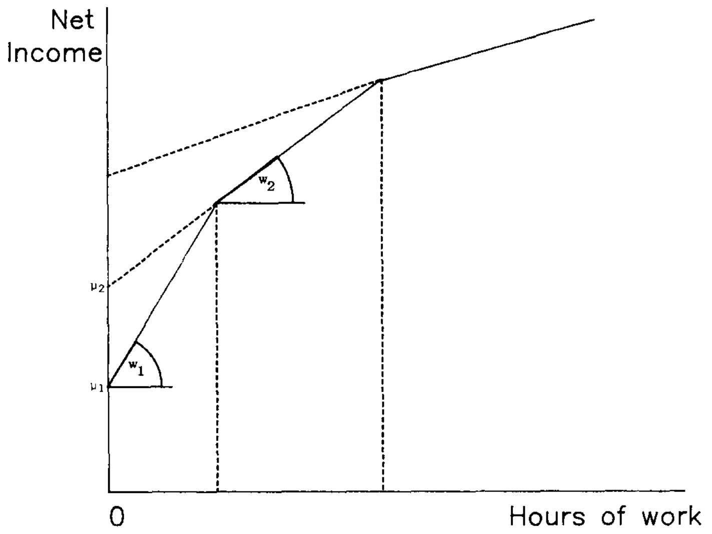

Fig. 1. Convex budget set.

where $\beta = \beta_f + \beta_{ff} \ln w - e$ and where w refers to the real wage. The LES contains only the first term on the right-hand side and hence labour supply is either monotonically increasing (if $\mu > \alpha_q$) or decreasing (if $\mu < \alpha_q$) in the wage. The GPF case contains the first and last terms and, given the estimated β_{ff} parameter values, we find that $\partial h/\partial w$ quickly becomes negative as w increases.

2.2. Simulation methodology

As an introduction to our methodology, consider some tax reform that results in the convex budget set illustrated in fig. 1. Such a post-reform constraint would arise, for example, with progressive income taxation. In this example the solution to the utility-maximisation problem can be found by exploiting the observation that a strictly quasi-concave utility function implies that the optimum is unique given the maximisation of the function on a closed and convex budget set. The procedure we adopt to solve for this unique optimum has been used by Hausman (1979) and Blomquist (1983). First predict desired labour supply under the first segment of the constraint, i.e. $h^* = h(w_1, \mu_1)$. If $h^* < 0$, then the optimum solution for the individual is to be a nonparticipant. If $0 < h^* < h_1$, we have found the optimum as that level of desired hours. If desired hours are greater than h_1, we move on to the next segment of the constraint. Here the marginal wage is w_2 and the corresponding intercept income is μ_2. If $h^* = h(w_2, \mu_2) < h_1$, then the optimum is at the kink h_1. If $h(w_2, \mu_2)$ is greater than h_2, we again move on to the next segment, and so on. In practice the algorithm is accelerated by starting the procedure at the pre-reform level of hours of work. This is likely to result in faster convergence to the solution since typically tax reforms do not result in drastic changes in the budget constraint and since labour supply may, in practice, not to be so very sensitive to changes in the constraint. A revealed preference argument can be used to show that starting at the pre-reform hours (or indeed at any level of hours) the first step the algorithm takes will be in the direction of the (global) optimum and convergence follows from the above discussion.

In the more general case of a nonconvex and even discontinuous constraint a unique optimum no longer necessarily exists. In such cases simulation can capitalize on the ease with which a solution can be found in the convex case by proceeding in the following fashion. First, identify all convex subsets of the nonconvex constraint; then apply the previous methodology to each convex subset computing utility for each (feasible) local solution (i.e. for each subset); and finally, choose the maximum maximorum. This is, in effect equivalent to considering the piecewise linear budget constraint as the union of a finite number of convex sets. Indexing these convex subsets by $j=1,\dots,J$, the overall indirect utility is given by

$$V^* = \max\{V_1(w_1,\mu_1),\dots,V_J(w_J,\mu_J)\}. \tag{7}$$

A potential problem that arises here is that the computation of the V_j's requires a knowledge of the support wage and intercept income pair. When a local optimum occurs at a kink in the constraint the supporting wage is not directly observed. For example, suppose the constraint illustrated in fig. 1 were a convex subset of some larger constraint set that were overall nonconvex, and suppose that the local optimum was found to be at the kink at h_1 hours. The computation of V_1 would require either a knowledge of the w, μ pair that would support h_1 as a tangency solution, i.e. w^*, μ^*, or a closed form for the corresponding direct utility function. As discussed in Meghir (1985) there is an exact correspondence between this problem and the analysis of rationing in Neary and Roberts (1980). The support or virtual wage corresponding to some ration level or kink in the budget constraint is the solution to $h(w^*,\mu_1+(w_1-w^*)h_1)=h_1$, where h_1 is the hours at the ration or kink, and μ^* has been replaced by $\mu_1+(w_1-w^*)h_1$.

Providing the compensated wage derivative of the labour supply curve is negative in a neighborhood of h_1 a unique virtual wage will exist. For example, with the LES model an explicit solution for w^* can be obtained by inverting the labour supply equation to yield:

$$w^* = \frac{\beta_1}{(1-\beta_f)}\,\frac{(\mu_1+w_1h-\alpha_q)}{(\alpha_f-h_1)}. \tag{8}$$

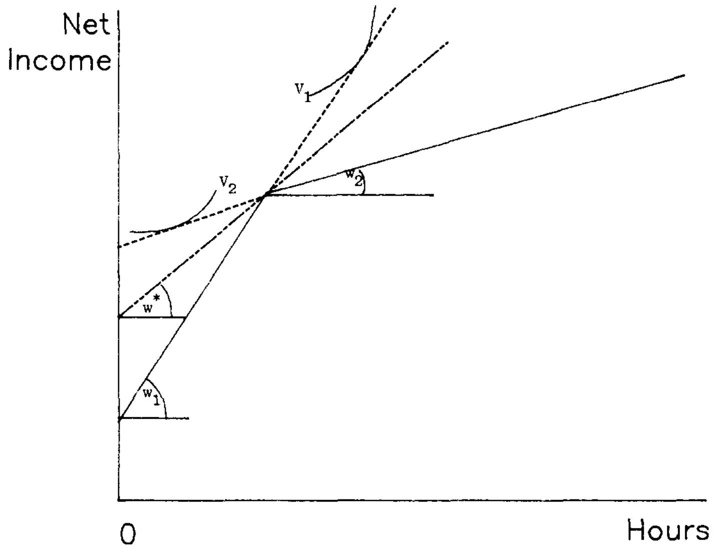

Fig. 2. Virtual wage.

However, the GPF case described above allows no such closed form solution and numerical methods need to be used. The foundation for such a method is illustrated in fig. 2. Here the unknown virtual wage is bounded above by the marginal wage w_1 and bounded below by w_2. The virtual wage w^* can be defined as the solution to $V(w^*, \mu^*) = \min_{w,\mu}\{V(w,\mu)|\mu = \mu_1 + h_1(w_1 - w)\}$. That is w^* is the slope of the line passing through the kink, which *minimises* the level of indirect (i.e. maximum) utility. Thus, as w is reduced from w_1 towards w_2 the level of utility falls from V_1 to some minimum at w^* and then rises to V_2. A numerical grid search routine can therefore exploit this property of w^* in its stopping rule, since $V(w,\mu)$ is strictly quasi-concave. In practice, such a calculation may be required for each individual in a large sample and it is worthwhile exploiting the bounds placed on w^* by the budget constraint. When the constraint contains discontinuities (where a small change in hours results in a large increase or decrease in net income) either the upper bound is infinite or the lower bound is zero. If, on the other hand, the constraint contains a 'spike' (such as would occur with fixed costs of work) the range of w^* is unbounded.

In the context of the labour supply of married women a further problem arises from the fact that a large proportion of women in any such sample will be nonparticipants. In this case the gross wage that they would face if they did work and the (estimated) stochastic error term of their preferences are both unknown. The error term (for both participants and nonparticipants) reflects taste heterogeneity and plays an important role in simulation as it appears in the marginal value of time, β. Consequently taste hetero-

Table 1

Proportion violating integrability.

Errors	Model	$h>0$	$h=0$
Included	LES	0.007	0.300
	GPF	0.355	0.451
Excluded	LES	0.000	0.000
	GPF	0.064	0.000

geneity is reflected in the virtual wages w^* and in general in all behavioural responses. Once generated these error terms remain constant across reforms, but are, of course, model-specific. For participants these problems do not occur since the wage is observed and the error term is taken to be the estimated residual. For nonparticipants the stochastic component of preferences is drawn from the truncated tail of the estimated normal distribution. This procedure ensures that the pre-reform nonparticipants are predicted to have nonpositive desired labour supply at their net wage rates. The net wages are generated by applying the tax/benefit rules using the model explained in Davies and Dilnot (1985) to a gross wage predicted from an auxiliary wage equation estimated over participants alone.[2]

2.3. Integrability conditions

For many forms of preferences the conditions under which the estimated labour supply function represents a solution to the utility-maximisation problem may not be satisfied at all data points. For example, in the linear labour supply equation the integrability condition is only guaranteed if the wage coefficient is positive and the income coefficient is negative.

As a guide to the severity of this problem for the labour supply equations used here the proportions of participants and nonparticipants which have incorrectly signed Slutsky substitution effects at their observed hours of work are given in table 1. Notice that setting the error terms e to zero typically results in integrability being satisfied at a larger proportion of sample points. The intuition behind this is that, as the estimated variance of the disturbance term is typically quite large, there is a significant probability that individuals will have errors large enough to take them outside the range of parameter values that would satisfy integrability. We expand on this point below. The calculations for participants on the left-hand side of table 1 use the estimated error terms. For nonparticipants there are no such estimates and, as described

[2]This procedure implicitly treats all nonworking women as voluntary nonparticipants. In contrast, the estimation method described in Blundell and Meghir (1986) is applied to data on workers alone and does not therefore make such an assumption.

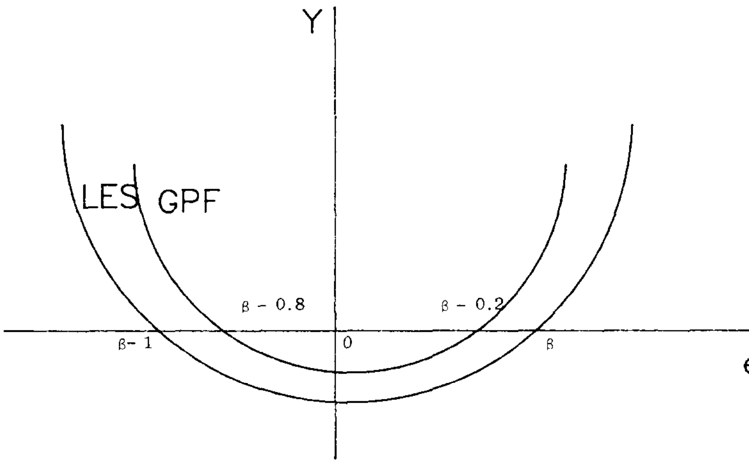

Fig. 3. Errors and integrability.

earlier, we generate the errors for these individuals by drawing from the truncated tail of the estimated truncated distribution for participants.

These observations can be better explained by investigating the requirements that the integrability conditions impose on the range in which the stochastic component of preferences can lie, compared with the range imposed by individual employment status. Integrability holds if the expenditure function is concave in the wage. That is, we require $\partial^2 E/\partial w^2 < 0$ where, corresponding to (1) the expenditure function is given by $E(w, U) = b(\cdot)V - a(\cdot)$. In both models considered in this paper $\partial^2 a/\partial w^2 = 0$ so that $\partial^2 E/\partial w^2 = V \cdot \partial^2 b/\partial w^2$. From (3), we have $\partial^2 b/\partial w^2 = b(\beta_{ff} + (\beta - e)(\beta - e - 1))/w^2$, where $\beta = \beta_f + \beta_{ff} \ln w$ and since b and w^2 are both positive, concavity of $E(\cdot)$ requires $(\beta - e)(\beta - e - 1) < -\beta_{ff}$.

It is instructive to write this condition as the following quadratic in e:

$$e^2 - e(2\beta - 1) + (\beta_{ff} + \beta(\beta - 1)) < 0. \tag{9}$$

Thus, concavity of $E(\cdot)$ is satisfied for values of e in the range $(\frac{1}{2})(2\beta - 1 \pm \sqrt{(1 - 4\beta_{ff})})$. Since β_{ff} is estimated to be 0.16 (see Appendix A for estimates) this range is from $\beta - 0.8$ to $\beta - 0.2$, i.e. a range of 0.6. Notice that for LES $\beta_{ff} = 0$ and the range of admissible e's ranges from β to $\beta - 1$, a range of 1.0. Thus, it is quite conceivable that while LES might be a less well fitting model than GPF so that the variance of the e's is higher, the range of admissible e's is so much larger that fewer individuals fail the condition. A diagrammatic interpretation to the integrability condition is facilitated by letting the expression in (9) equal Y and plotting e against Y. This yields fig. 3, where the admissible range is the horizontal axis between the intercepts. In the GPF case the centre point of this range depends on w since $\beta=\beta_f+\beta_{ff}\ln w$. For sufficiently small or high w the GPF range of e's may not include zero. In the GPF case, therefore, excluding the error does not necessarily guarantee concavity, although in practice the model seems to be relatively successful in this respect. In contrast the estimated LES case with error terms set to zero is always consistent with a concave expenditure function. This occurs because the estimated β always falls between 0 and 1 and supernumerary income is positive for all sample points. However, ignoring the stochastic component of preferences in simulation would not respect the observed employment status of individuals in the data.

Whenever the use of estimated errors implies a lack of consistency with the underlying economic theory some difficulties will arise during simulation. From the discussion above the problem is obviously more serious in the context of the more flexible GPF model. It would be inappropriate to simply exclude those observations that exhibit inconsistency with the theory since the result would be an unrepresentative sample. The simplest alternative strategy would be to assume that those individuals who fail integrability do not change their labour supply in response to a change in the constraint. That is, individuals whose behaviour the model cannot explain in a consistent fashion are assumed, for simulation purposes, to be rationed at their observed hours. Given that most concavity failures occur amongst non-participating women with young children, and since their reservation wage is likely to be large, this may be less of a restriction than it first appears. However, two other alternatives suggest themselves. First, one could attempt to impose integrability using the existing parameter estimates by choosing an e to satisfy integrability (subject to the restrictions imposed by the observed employment status, etc.) for those individuals who fail integrability at their estimated e. Secondly, one could attempt to impose integrability conditions at the estimation stage.

In the linear labour supply and the linear earnings cases the imposition of integrability is relatively easy to achieve [see Hausman (1981) and King (1987)] and although it is possible in other cases it may not always be advisable. As Diewert and Wales (1987) have shown, imposing integrability conditions on the parameters of commonly used flexible functional forms can considerably reduce their ability to approximate a wide range of behaviour. This problem is more serious the greater is the variation in wages and incomes over which the model is being estimated. Cross-section data sets clearly display wide variation, but allowing parameters to be dependent on individual and household characteristics, as outlined in eqs. (2), (3) and (5), mitigates against this problem in estimation. However, in simulation where we have to evaluate behaviour for the *same* individual under two possibly very different wage and income levels, the problem re-emerges. Thus, a stochastic specification that forces integrability on to the estimates may well lead to the estimates having features that are not necessarily reflections of behaviour in the data.

In the tax simulation results presented below integrability failures are dealt with using the simplest of the approaches described above restricting labour supply in these cases to be unaffected by the reform. For comparison, we also present in an appendix some results where the random preference error is chosen from the range which satisfies concavity. Since the imposition of integrability at the estimation stage suffers from the drawbacks mentioned above, we do not pursue this approach here. Nevertheless, this procedure seems worthy of investigation and is work we intend to pursue in future.

3. The reform of personal taxation

3.1. The reforms explained

The reforms we are concerned with here relate to the contrast between the tax treatment of the income of couples as opposed to that for single individuals. This problem has long been of interest in the United States where there have been significant changes in the tax treatment of married couples over the last 40 years [see Bittker (1975) for some history and Feenberg (1982) for some analysis of the two-earner couple deduction introduced in the United States in 1981].

The tax, benefit, and compulsory national insurance (social security) systems in the United Kingdom produce a complicated budget constraint with several significant nonconvexities and at least one discontinuity. The Housing Benefit (HB) and Family Income Supplement (FIS) means tested benefits give rise to significant nonconvexities, the national insurance contributions (NIC) system gives rise to both nonconvexities and a discontinuity (three since 1986), while the income tax system itself is characterised by increasing marginal rates as income increases. In the United States similar points apply: AFDC (the equivalent of SB), and FICA (the equivalent of NIC) result in nonconvexities while the income tax system alone is progressive overall.

The salient features of the present U.K. tax system can be explained relatively simply. Married men receive a tax allowance approximately 60 percent larger than the allowance to which single individuals and married women are entitled.[3] This special treatment for married men is now viewed as something of a historical anomaly and it is the reduction of the married man's allowance (MMA) that is the common feature of the two reforms considered below. Married individuals who earn more than their allowance face a marginal tax rate which is determined by the combined income of the

[3]The allowance for married women is against earned income only.

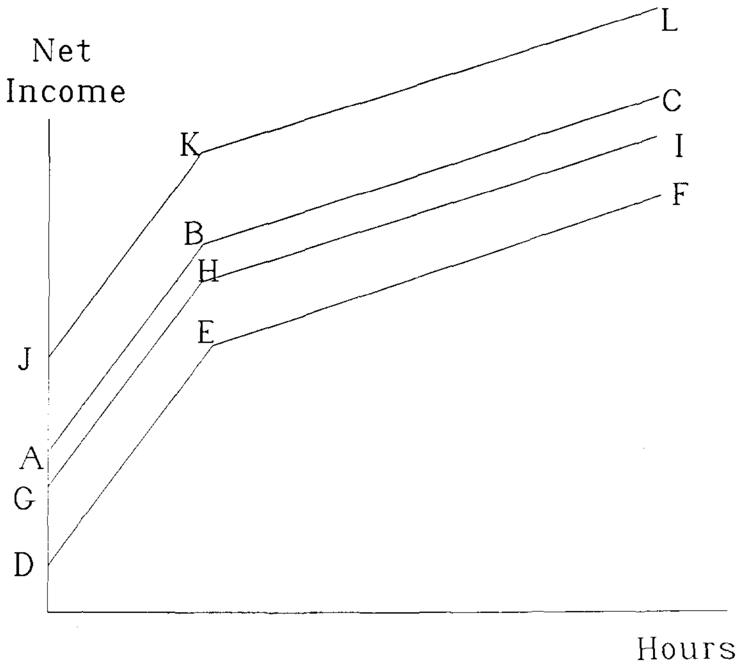

Fig. 4. CB reform constraints.

couple.[4] In practice more than 90 percent of married men face the standard tax rate since this covers a very large range of incomes.

Interest in the taxation of married couples in the United Kingdom stems from two government discussion documents [HMSO (1980, 1986)]. The first considered a range of options for the reform of the taxation of couples. The option favoured by most respondents to that document [see Kay and Sandler (1982)] was to reduce the MMA to the level of the single allowance (SA) and use the extra revenue to increase Child Benefit (CB: a lump-sum weekly payment per dependent child). We denote this reform as CB. The 1986 and 1980 discussion documents both proposed the reduction of the MMA to the level of an SA and that husband and wife be allowed to transfer their SA's to minimise the household's tax liability if they so wished. We refer to this as the TA (Transferable Allowances) reform. In both cases the aggregation of a couple's incomes is also abandoned. These reforms are particularly useful for comparing model specifications for female labour supply since one (CB) generates only income effects while the other (TA) generates some large changes in marginal tax rates faced by many married women.

The essential features of the CB reform is captured in fig. 4 which, for the

[4]There exists the opportunity for couples to opt for separate taxation and be treated as single individuals.

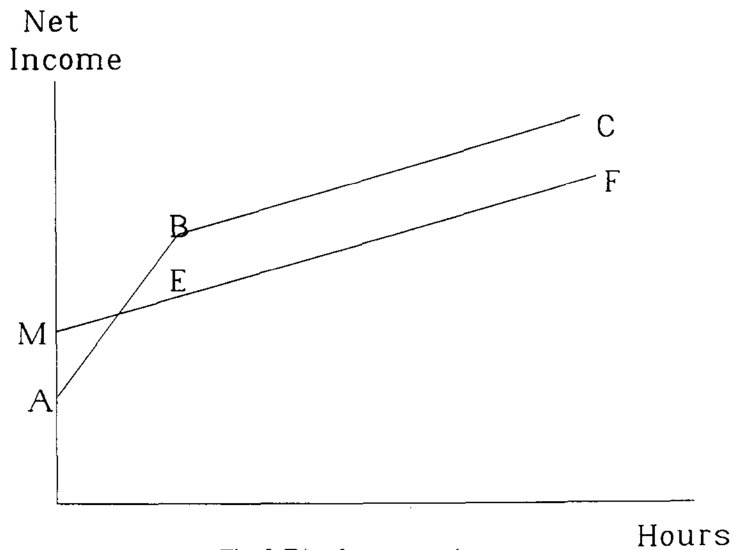

Fig. 5. TA reform constraints.

sake of diagrammatic simplicity, abstracts from the complications induced by the NIC, HB, and FIS systems, and by higher rate tax bands. The pre-reform constraint is *ABC*, where *B* occurs at the point where the wife's earnings exceed her allowance, SA. Above SA her income is taxed at a rate t which is determined by the sum of the couple's taxable incomes but which is generally the standard rate. The CB reform reduces the MMA to the level of the SA and uses the increase in revenue to increase CB to CB′ (and to abandon the aggregation rule). For a childless couple the constraint drops to *DEF*, where *AD* is t(MMA − SA). For a household with one child the constraint moves to *GHI*, where *AG* is $(\mathrm{CB}'-\mathrm{CD}^0)-\mathrm{AD}$ and $\mathrm{CB}'-\mathrm{CB}^0$ is the change in child benefit. Similarly, with two children the constraint becomes *JKL*, where *AJ* is $2(\mathrm{CB}'-\mathrm{CB}^0)-\mathrm{AD}$. As can be seen from fig. 4 the reform is unlikely to have much impact on the marginal tax rates faced by individuals. The only significant changes occur for women who pay higher rate tax under the present aggregation rule.

On the other hand the abolition of the MMA and the introduction of Transferable Allowances – the TA reform – results in the situation portrayed in fig. 5, at least for the vast majority of couples. As before we abstract, in the diagram, from the complications associated with higher rate taxpayers, separate election, means tested benefits, and national insurance since they do not affect the substance of the argument. As before, *ABC* again represents the pre-reform constraint. *MEF* represents the post-reform constraint. The TA reform allows a small increase in SA to ensure revenue neutrality. Thus, *EF* in fig. 5 lies slightly above *EF* in fig. 4. In fig. 5, if the wife supplies no work

Table 2

1981 tax parameters.

	MMA	SA	CB	Transferable	Aggregation
Actual	41.25	26.44	4.75	No	Yes
CB reform	26.44	26.44	8.25	No	No
TA reform	27.30	27.30	4.75	Yes	No

to the market the household is better off by AM since the husband can add his wife's allowance to his own SA and $t(2\text{SA}-\text{MMA})$ is positive since MMA is approximately 60 percent more than as SA. On the other hand, if the wife has earnings in excess of the SA the household is worse off by $t(\text{MMA}-\text{SA})$. If the wife earns less than SA she can either transfer her whole SA to her husband, and so the *household* pays the tax liability on the wife's earnings from those earnings, or she can transfer only the unused portion of her SA, so that the *household* pays the tax liability on the wife's earnings out of the husband's earnings. In either case the results will lie on the line ME in fig. 5 and in either case the tax rate the household faces on an extra £1 of wife's earnings is the same.

Notice that the TA reform increases the marginal tax rate faced by all women with earnings below the SA (provided the husband earns more than 2SA), including those earning zero. The TA reform also gives rise to significant changes in net incomes for couples who are currently both taxpayers. The most significant aspect of fig. 5 is that the TA reform while ostensibly reducing the MMA is in practice equivalent to an increase in the MMA and the abolition of the wife's SA. In table 2 we present a summary of the relevant parameters (in £ per week) of the pre-reform (April to October 1981) system and of the two reforms.

The sample used in simulation is 1290 households drawn from the 1981 U.K. Family Expenditure Survey. Details of the data are provided in appendix B. The data relate only to the period from April to October 1981 in an effort to ensure that all households face the same set of tax/benefit parameters. The budget constraint is generated by a program that calculates net incomes and net wage rates on the basis of the tax/benefit rules and the data on gross wage rates, housing costs, number of children, etc. The actual net wages and incomes may differ from those calculated by the program because of nontake-up of benefits, nonpayment of taxes, and administrative delays in awarding/withdrawing benefits and/or allowances. Our solution to this problem is to simulate the behaviour of the sample on the calculated budget constraint of the pre-reform tax system. Since this simulated pre-reform position depends on the model used we have, in fact, two pre-reform data sets; one corresponding to each labour supply model. Appendix B gives

Table 3

CB results by numbers of children.

	% of sample	Δh	$\Delta y\|_h$	Δy	EG
GPF					
0	34.1	0.64	−4.13	−3.27	−4.13
1	21.8	0.14	−0.89	−0.67	−0.85
2	32.1	−0.29	2.48	2.06	2.51
3+	11.9	−0.55	6.56	5.89	6.57
	100	0.09	−0.02	0.10	−0.01
LES					
0	34.1	0.68	−4.13	−3.09	−4.08
1	21.8	0.29	−0.89	−0.44	−0.86
2	32.1	−0.31	2.48	2.07	2.50
3+	11.9	−0.73	6.62	5.92	6.69
	100	0.10	−0.01	−0.015	0.02

details of both data sets and while the two differ little the GPF data set contains one fewer individuals who was predicted to have pre-reform hours in excess of 90 per week and one more nonparticipant who is observed in reality to be working and predicted to be working under the LES model but predicted to be not working with the GPF model.

The gains, losses, and efficiency effects, etc. of those reforms described in the next two subsections were computed using our microcomputer policy simulation package called SPAIN [see Symons and Walker (1986) for further details]. This program accounts for all aspects of the tax/benefit system including those which were ignored for the sake of expositional clarity in figs 4 and 5.[5]

3.2. CB reform simulation results

Table 3 gives a breakdown of the main effect by number of children for the two models. The pattern of hours changes are almost entirely the result of income effects. $\Delta y|_h$ shows the net income changes that would occur if there were no behavioural responses (i.e. given hours) and these differ only through the models having slightly different pre-reform data sets. Δh shows the hours changes. Δy shows the net income changes given the change in hours. Notice that $\Delta y > \Delta y|_h$ if $\Delta h > 0$ and vice versa, EG is the change in equivalent income arising from the reform. This is a money metric of welfare change due to King (1983) and is bounded below by $\Delta y|_h$ when there is no change in marginal wage rates (or when there is a zero compensated substitution effect if leisure is a normal good). Since marginal wage rates

[5]We assume all couples take up all eligible means tested benefits and opt for separate taxation whenever this would be beneficial.

Table 4

CB reform labour supply changes.

	NP	1–8	9–16	17–24	25–30	31–37	38–45	46+	Post reform dist.
					GPF model				
NP	518	3							521
1–8	6	65	3						74
9–16		4	114	3					121
17–24			7	141	8				156
25–30				5	62	1			68
31–37					10	122	1		133
38–45						23	186		209
46+								7	7
Total	524	72	124	149	80	146	187	7	1289
					LES model				
NP	502	4							507
1–8	21	61	3						85
9–16		5	110	5					120
17–24		1	9	138	8				156
25–30				9	59	1			69
31–37					13	125	1		139
38–45						20	187		207
46+								7	7
Total	523	71	122	153	80	146	188	7	1290

change for only a minority of the sample the *EG*'s are very close to the $\Delta y|_h$ values.

Differences across models in table 3 are clearly minor. The reason for this insensitivity is that the reform gives rise predominantly to income effects and hence the only model differences that come into play are the ones reflected in the marginal budget shares. Essentially these are similar for individuals across models because the random preference error component in the LES marginal budget share adjusts to compensate for the absence of $\ln w$ from this specification. Table 4 gives a broader view of the labour supply responses. This highlights the only significant difference across the models – the difference in the number of new participants. In the LES model 21 non-participants start work, while with the GPF model there are only six such women. This seems to be due to the greater sensitivity of the reservation wage to reductions in μ induced by the reform in the LES model compared with the GPF. In other words, the reduction in μ for childless nonparticipants in the LES model induces more substantial decreases in reservation wages and hence a larger number of new participants.

Notice from table 3 that there is some labour supply response in aggregate, albeit a small one. This arises because the gainers are concentrated in a group that are less likely to work than the losers. Hence, the increases in hours of the losers is not matched by a decrease in hours by the gainers

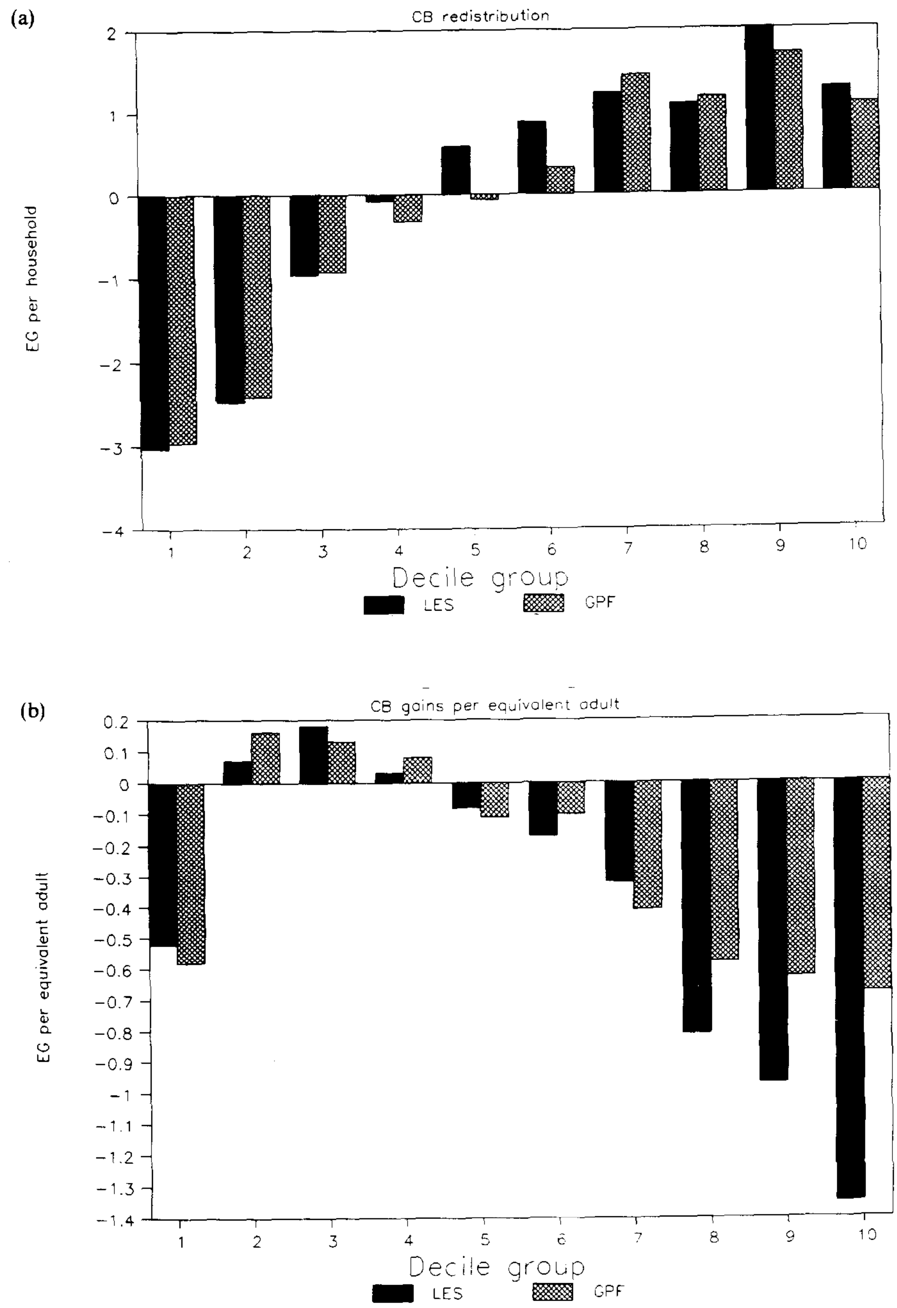
(a)
CB redistribution
EG per household
2
1
0
−1
−2
−3
−4
1 2 3 4 5 6 7 8 9 10
Decile group
LES
GPF
(b)
CB gains per equivalent adult
EG per equivalent adult
0.2
0.1
0
−0.1
−0.2
−0.3
−0.4
−0.5
−0.6
−0.7
−0.8
−0.9
−1
−1.1
−1.2
−1.3
−1.4
1 2 3 4 5 6 7 8 9 10
Decile group
LES
GPF

Fig. 6

Table 5

CB revenue and deadweight loss effect.

Model	Change in government revenue per household	Change in DWL per household
GPF	0.13	−0.14
LES	0.17	−0.19

because a larger proportion of the latter are already nonparticipants. Of course, these simulated responses may depend on the method used for dealing with integrability failures described in subsection 2.3 above. In appendix C we present results that are obtained from the alternative practice of adjusting the error term of concavity violators to ensure integrability. The results there suggest that our procedure is comparable with this alternative.

The consequences of the CB reform tax revenue and deadweight loss in £ per household per week are spelled out in table 5. The government revenue rises more with LES than GPF since the former implies a slightly larger hours increase. Finally, fig. 6(a) shows the effect of the reform on the distribution of equivalent incomes. The bars shows the the average EG for each decile of the preform utility distribution. The overall skewed shape arises because children are concentrated in the centre of the equivalent income distribution and above. This is why in fig. 6(a) the poor appear to lose.

Although utility adjusts for the way household composition affects the allocation between time and goods, utility measures derived from revealed behaviour can only ever identify household welfare up to some monotonic transformation which may well depend on household characteristics. Econometric analysis cannot, therefore, completely determine the correct welfare function with which to rank households. If per capita [using the Supplementary Benefit (SB) scales] utility is used to rank households, children now tend to appear in the 'poor' households and so the reform redistributes toward the lower deciles. These results are presented in fig. 6(b). Fig. 6(a) implicitly treats household welfare as a local public good in so far as it attributes the estimated level of welfare V to each household independently of household size. In other words, two households which may differ in household composition and other determinants of V but for which indirect utility is equal are treated equally. In contrast, fig. 6(b) treats household welfare a private good and we divide that across household members using the SB scales. This results in some households appearing in different deciles in the two figures.

3.3. *TA reform results*

Table 6(a) gives a breakdown of the main effects of the reform by type of

Table 6

(a) TA results for the GPF model.

	Sample (%)	$\Delta y\|_h$	Δt	Δh	$\Delta h\|_u$	Δy	EG
Nonparticipants	0.41	4.10	0.30	0	–	4.10	4.10
Earnings <£12.7	0.06	1.31	0.30	−2.59	−1.76	−1.16	1.72
£12.7–£26.44	0.09	−2.02	0.30	−1.09	−1.68	−3.14	−1.68
Basic rate	0.40	−3.93	0.01	0.74	−0.01	−2.97	−3.93
Worse off, high	0.04	−2.83	−0.08	0.64	0.16	−2.01	−2.80
Better off, high	0.01	8.26	−0.14	−0.16	0.18	7.78	8.33
All participants	0.59	−2.81	0.06	0.12	−0.41	−2.52	−2.74
All households	1.00	0.01	0.16	0.07	–	0.18	0.05

(b) TA results for the LES model.

	Sample (%)	$\Delta y\|_h$	Δt	Δh	$\Delta h\|_u$	Δy	EG
Nonparticipants	0.41	4.10	0.30	0	–	4.10	4.10
Earnings <£12.7	0.06	1.30	0.29	−5.73	−2.01	−3.83	2.58
£12.7–£26.44	0.09	−2.06	0.31	−5.03	−5.11	−7.18	−0.86
Basic rate	0.40	−3.93	0.00	0.75	−0.02	−2.94	−3.93
Worse off, high	0.03	−2.71	−0.08	1.47	1.03	−0.31	−2.62
Better off, high	0.02	7.24	−0.10	−0.36	0.59	8.74	7.49
All participants	0.59	−2.80	0.07	0.69	−0.87	−3.15	−2.46
All households	1.00	0.00	0.16	−0.41	−0.52	0.21	0.20

household for the GPF model. We have distinguished between six different sorts of household. Fig. 7(a) illustrates the first four types and fig. 7(b) illustrates the remainder. In fig. 7(a) nonparticipants clearly remain nonparticipants (move from a^0 to a') since they face higher incomes and lower net wage rates. The income gain observed in table 6(a) is £4.10 on average made up of £4.03 gain for households with standard rate (30 percent) taxpayer husbands ($(2SA'-MMA)t=4.03$ for $t=0.3$) and more for the relatively few higher rate taxpayers. Basic rate taxpayers move from b^0 to b' in fig. 7(a) and incur only income effects induced by the reduction in μ of £3.93 ($=0.3(MMA+SA-2SA')$). Since Δt is small, EG is approximately $\Delta y|_h$. In fact a small proportion of these households experience an EG less than $\Delta y|_h$ because their behavioural change takes them across the National Insurance floor which represents a discontinuity in the constraint and a change in marginal rate. Notice that $\Delta y > \Delta y|_h$ since the income effect on hours induces an increase in hours of 0.74 which induces a reduction in the loss of income.

For those households where the wife earns less than £12.70 per week the reform induces a move from, say, c^0 to c' in fig. 7(a). The rise in income of £1.21 in table 6(a) and increased marginal rate both induce a reduction in hours of 2.59 of which 1.76 hours are accounted for by the substitution effect alone. The hours reduction is so large that income *falls* by £1.16 post-response. For those wives earning between £12.70 and £26.44 per week the reform would leave them on the standard tax rate and £2.02 worse off. The

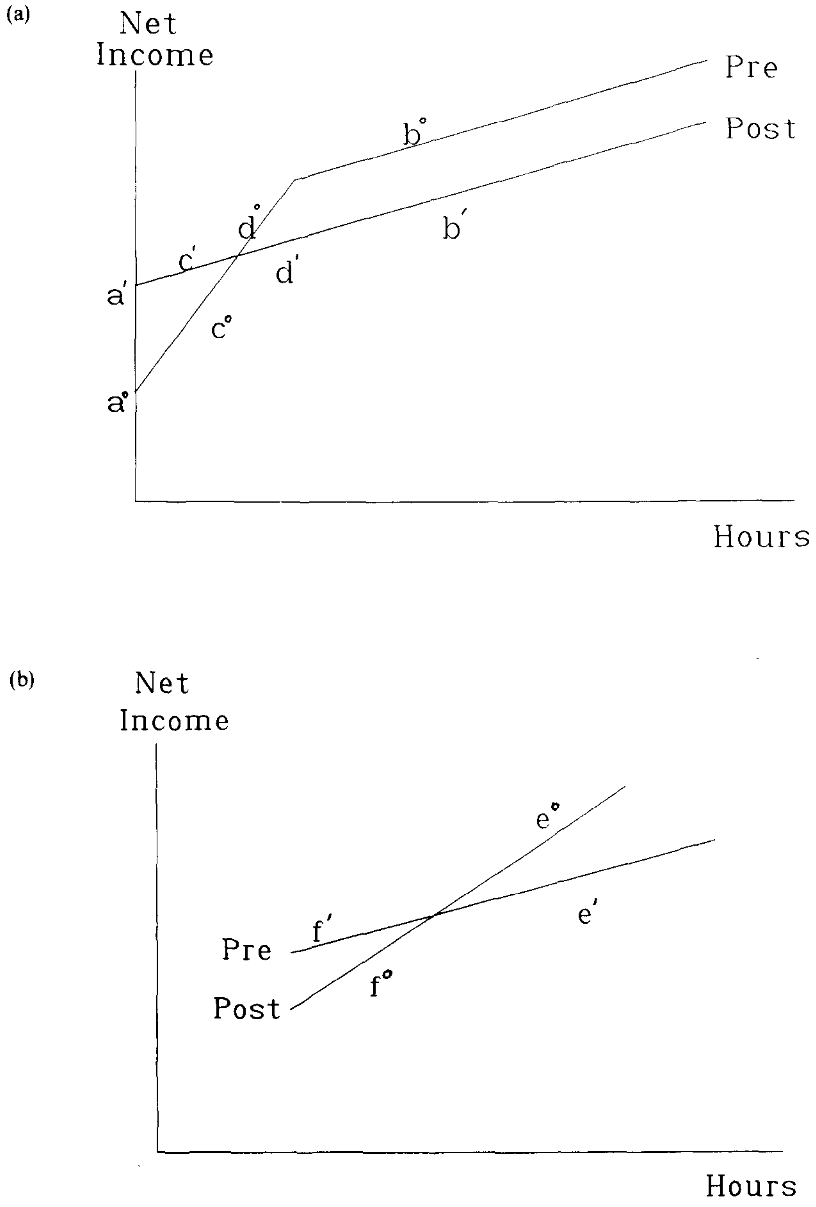

Fig. 7. (a) TA reform effects; (b) TA reform – higher rate taxpayers.

substitution effect is large at -1.68 and this dominates the positive income effect so that the gross effect on hours is a reduction of 1.09. Thus, income falls post-response by £3.14 as opposed to £2.02 pre-response.

Fig. 7(b) shows the two possible household types that are subject to higher rate tax. Most households containing higher rate tax paying women are worse off since the increased tax bill of the husband due to the reduction in the MMA more than offsets the reduced two bill of the wife that arises because she now pays tax at the standard rate rather than at her husband's rate. Thus, such households face a lower tax rate on the wife's earnings which induces a substitution effect of 0.16 hours reinforced by the income effect arising from being worse off to give a gross hours effect of 0.64 hours increase. Thus, such households are, post-reponse, worse off by £2.01 rather than £2.83 pre-response. Since there is a small decrease in marginal tax rate the money metric *EG* is $-£2.80$, slightly larger than $-£2.83$. These households are depicted in fig. 7(b) by the move from d^0 to d'. There are, however, a few households where the reduced tax bill on the wife's earnings is not offset by the increased tax bill of the husband. Such cases occur when the wife's earnings are relatively large and total household income is not too much greater than the higher rate threshold. This case is depicted by the move from e^0 to e' in fig. 7(b). Here income and substitution effects counteract each other so that the 0.18 hour compensated effect is outweighed by the income effect to give a 0.16 hour *decrease* gross. The *EG* is greater than $\Delta y|_h$, again because of the reduction in marginal rate.

The results for the LES model are presented in table 6(b) and are qualitatively the same as for the GPF model. However, the LES model tends to produce larger responses to both wage and income changes. These differences are highlighted in tables 7(a) and (b) where the labour supply effects are documented. The LES model implies a significant reduction in participation while this effect is relatively modest in the GPF model.

Table 8 illustrates the revenue and deadweight loss effects of the TA reform. On efficiency grounds both models indicate that TA is preferable to the status quo. Fig. 8(a) gives the breakdown of *EG* by decile group of the pre-reform equivalent income distribution. The skewed shape reflects the fact that it is low-wage, two-earner couples that lie at the bottom of the household welfare distribution and higher rate taxpayers are concentrated towards the top of the distribution. Fig. 8(b) provides results when households are ranked by per capita equivalent income [see the discussion of fig. 6(b) above] and, unlike the results for the CB reform, show very few changes in comparison with fig. 8(a).

Finally, in addition to the worries concerning the scaling of indirect utility or equivalent income by some measure of household size, one should also be careful in drawing welfare conclusions that do not consider intertemporal effects of reforms. Clearly, where gains and losses are determined by the

Table 7

(a) TA labour supply effects: GPF model.

	NP	1–8	9–16	17–24	25–30	31–37	38–45	46+	Post reform dist.
NP	523	8							532
1–8		62	19						81
9–16		1	100						101
17–24			5	139					444
25–30				10	57				67
31–37					23	124			147
38–45						21	185		206
46+								7	7
Total	523	71	124	149	80	145	185	7	1285

(b) TA labour supply effects: LES model.

	NP	1–8	9–16	17–24	25–30	31–37	38–45	46+	Post reform dist.
NP	523	46	6	1					576
1–8		25	55	1					81
9–16			52	13					65
17–24			9	122	4				135
25–30				16	50				66
31–37					26	123			149
38–45						23	187		210
46+							1	7	8
Total	523	71	122	153	80	146	188	7	1290

Table 8

Revenue and deadweight loss effect of TA.

Model	Change in government revenue per household	Change in DWL per household
GPF	0.17	−0.22
LES	0.01	−0.19

presence of children or labour market status of the wife there may be considerable variations in gains and losses across households' lifetimes.

4. Conclusions

Here we have outlined a methodology for evaluating the implications of tax reforms on the labour supply behaviour and welfare of a large sample of households. The procedure is embodied in a microcomputer program and is sufficiently general to simulate the effects of a wide variety of tax/benefit reforms allowing for all of the interplay between taxes, social security

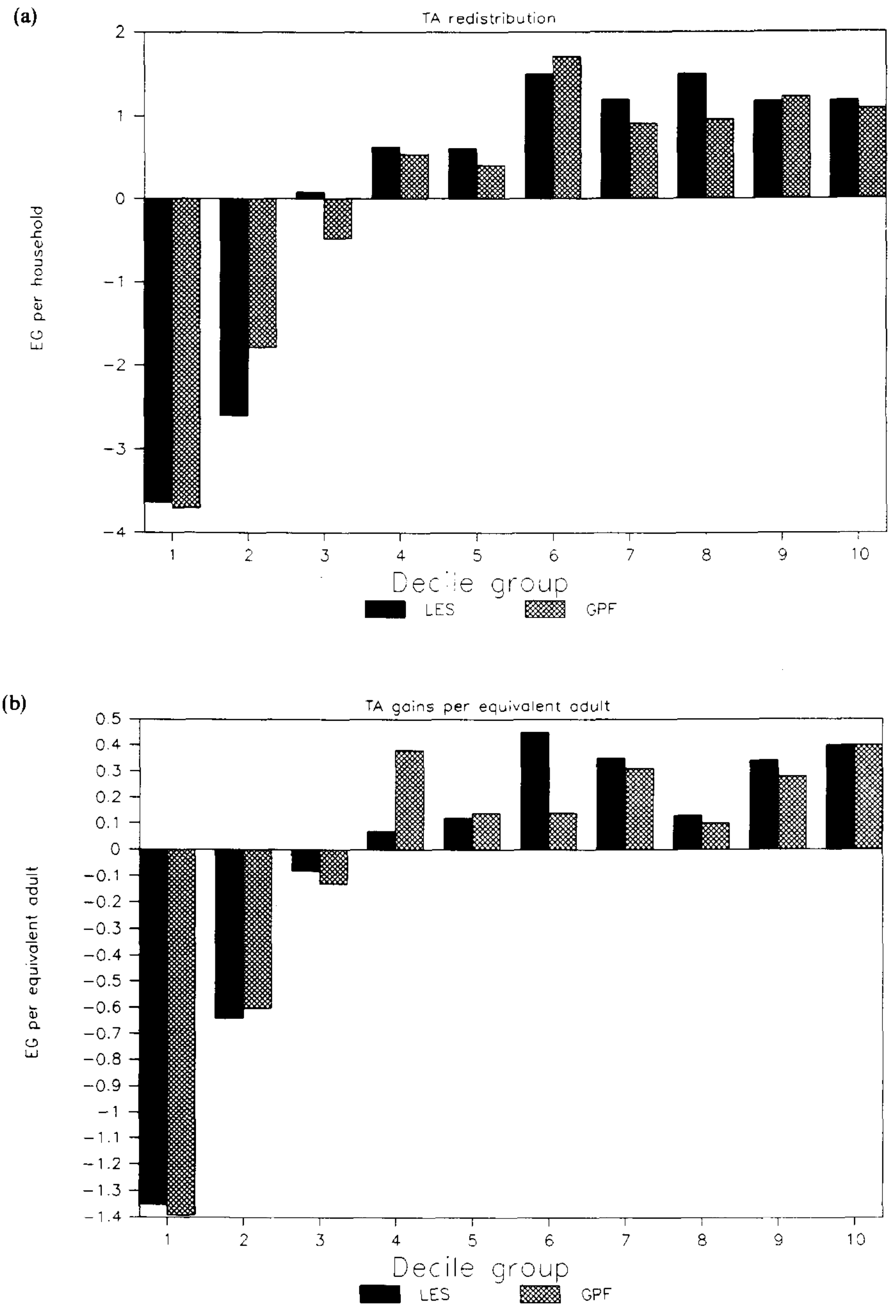
(a)
TA redistribution
EG per household
2
1
0
-1
-2
-3
-4
1 2 3 4 5 6 7 8 9 10
Decile group
LES
GPF
(b)
TA gains per equivalent adult
EG per equivalent adult
0.5
0.4
0.3
0.2
0.1
0
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
-0.7
-0.8
-0.9
-1
-1.1
-1.2
-1.3
-1.4
1 2 3 4 5 6 7 8 9 10
Decile group
LES
GPF

Fig. 8

contributions, and means tested benefits. We have used this simulation and analysis routine in this paper to investigate the extent to which alternative specifications for the labour supply equation may generate different conclusions. Our analysis focused on the labour supply of married women since this is the group that appears most likely to respond to wage and income changes. We therefore considered two alternative specifications for the labour supply of married women. The major difference between these was that one, the LES model, restricts the labour supply curve to be monotonic in the wage. This restriction is a common one in the empirical labour supply literature. The alternative and more data coherent model allows for both backward bending and forward sloping labour supply to occur simultaneously.

The labour supply models were used to investigate the effect of reforms to the taxation of family income. This aspect of the tax system is one which has preoccupied policymakers in many countries yet little work has been conducted on the likely effects of reforms. The specification of the two reforms are drawn from recent proposals for the U.K. tax system. One involves the abolition of the married man's allowance, the abolition of the aggregation of a couple's incomes for tax purposes, and an increase in lump-sum child benefit. In the United Kingdom, where the great majority of couples face the standard tax rate, this reform has predominantly income effects. The second reform that we consider also abolished the married man's allowance and income aggregation but uses the revenue to allow couples to transfer their allowances to minimise the household's tax liability. This reform effectively increases the marginal tax rate on nontaxpaying wives by 30 percent and also has income effects.

The choice of model specification has practically no effect on the behavioural implications of the first reform except where marginal tax rate changes are involved, and for nonparticipants. In aggregate these differences are minor except for government revenue which is higher for the LES model because of its larger incentive effects to higher rate taxpayers. In the second reform the large marginal tax rate changes give rise to very significant differences in behaviour. Moreover, in the comparative exercise the choice of model was found to make a difference to policymakers ranking of reforms. It would appear that not only are simple labour supply equations liable to impose theoretically unattractive and empirically rejectable restrictions, but it is also possible that they mislead policy analysts when it comes to weighing up alternative routes to tax and benefit reforms.

Appendix A: Model estimates (standard errors)

	GPF		LES	
β_f^0	0.267	(0.018)	0.234	(0.020)
β_f'	0.253	(0.032)	0.567	(0.088)
β_f''	0.162	(0.021)	0.248	(0.040)
β_f'''	0.034	(0.018)	0.089	(0.045)
$\beta_f^a \times 100$	0.034	(0.006)	0.047	(0.007)
$\beta_f^{aa} \times 100$	0.016	(0.006)	0.011	(0.007)
β_f^e	−0.010	(0.002)	−0.003	(0.003)
β_{ff}	0.159	(0.013)	–	–
α_f^0	42.280	(0.710)	41.720	(0.530)
α_f'	–	–	17.250	(8.120)
α_f''	–	–	3.290	(2.560)
α_f'''	–	–	3.920	(2.570)
α_q^0	16.940	(1.540)	19.970	(0.990)
α_q'	0.620	(3.330)	8.900	(2.950)
α_q''	−0.510	(1.960)	2.020	(1.450)
α_q'''	3.010	(1.010)	2.960	(1.010)
σ_e	0.172	(0.007)	0.202	(0.007)
Log L	−5136.73		−5184.21	

Appendix B: Data

B.1. LES model

Sample size: Households with participating women: 767
Households with nonparticipating women: 523

Analysis for households with participating wives:

Variables	Mean	Standard deviation	Minimum	Maximum
Female hours	26.30	12.03	1	80
Male hours	39.87	5.84	24	96
Children 0–4	0.16	0.44	0	4
Children 5–10	0.44	0.71	0	3
Children 11–18	0.44	0.75	0	3
Female wage	2.20	1.09	0.5	11.54
Male wage	3.68	1.46	0.65	16.31
Income (μ)	38.33	39.34	−54.43	247.96
Female age	36.06	10.10	17	59
Education	16.10	2.39	0	25

Analysis for households with nonparticipating wives:

Variables	Mean	Standard deviation	Minimum	Maximum
Female hours	0	0	0	0
Male hours	40.51	6.89	16	110
Children 0–4	0.76	0.81	0	3
Children 5–10	0.50	0.73	0	3
Children 11–18	0.32	0.68	0	4
Female wage	1.91	1.53	0.5	4.03
Male wage	3.98	1.84	0.68	17.31
Income (μ)	71.69	38.32	20.92	398.93
Female age	34.87	10.30	19	59
Female education	16.00	2.24	0	25

B.2. GPF model

Sample size: Households with participating women: 765
Households with nonparticipating women: 524

Analysis for households with participating wives:

Variables	Mean	Standard deviation	Minimum	Maximum
Female hours	26.53	12.05	1	80
Male hours	39.86	5.84	24	96
Children 0–4	0.16	0.44	0	4
Children 5–10	0.43	0.70	0	3
Children 11–18	0.44	0.75	0	3
Female wage	2.21	1.09	0.5	11.54
Male wage	36.73	1.44	0.65	16.31
Income (μ)	37.78	39.05	−54.43	247.96
Female age	36.05	10.11	17	59
Female Education	16.09	2.38	0	25

Analysis for households with nonparticipating wives:

Variables	Mean	Standard deviation	Minimum	Maximum
Female hours	0	0	0	0
Male hours	40.52	6.82	16	110
Children 0–4	0.76	0.81	0	3
Children 5–10	0.49	0.73	0	3
Children 11–18	0.32	0.68	0	4
Female wage	1.91	0.53	0.5	4.03
Male wage	3.97	1.84	0.68	17.31
Income (μ)	71.69	38.28	20.92	387.93
Female age	34.89	10.30	19	59
Female education	16.00	2.34	0	25

Appendix C: Integrability and error distribution effects on hours – TA reform (results for the CB reform are identical in table 4)

C.1. LES model

C.1. LES model

Post-reform hours	Pre-reform hours						
	NP	1–8	9–16	17–24	25–30	31–37	38+
NP	489	45	3	0	1	0	0
1–8	0	29	59	1	0	0	0
9–16	0	0	55	13	0	0	0
17–24	0	0	6	124	4	0	0
25–30	0	0	0	12	54	0	0
31–37	0	0	0	0	23	128	1
38+	0	0	0	0	0	21	190

C.2. GPF model

Post-reform hours	Pre-reform hours						
	NP	1–8	9–16	17–24	25–30	31–37	38+
NP	524	7	0	0	0	0	0
1–8	0	66	19	0	0	0	0
9–16	0	1	114	0	0	0	0
17–24	0	0	13	203	0	0	0
25–30	0	0	0	20	130	0	0
31–37	0	0	0	0	25	130	0
38+	0	0	0	0	0	18	53

References

Arrufat, J.L. and A. Zabalza, 1986, Female labour supply with taxation, random preferences and optimization errors, Econometrica 54, 47–64.

Ashenfelter, O. and J. Ham, 1979, Education, unemployment and earnings, The Journal of Political Economy 87, S99–S166.

Bittker, B., 1975, Federal income taxation and the family, Stanford Law Review 27, 1392–1467.

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.W. and C. Meghir, 1986, Selection criteria for a microeconometric model of labour supply, Journal of Applied Econometrics 1, 55–81.

Blundell, R.W. and I. Walker, 1986, A life-cycle consistent empirical model of family labour supply using cross-section data, Review of Economic Studies 53, 539–558.

Burtless, G. and J.A. Hausman, 1978, The effect of taxation on labour supply: Evaluating the Gary negative income tax experiment, Journal of Political Economy 86, 1103–1130.

Davies, E. and A.W. Dilnot, 1985, The IFS tax and benefit model, Institute for Fiscal Studies Working paper, 58.

Diewert, W.E., 1974, Applications of duality theory, in: M.D. Intriligator and J.W. Kendrick, eds., Frontiers of quantitative economics, vol. II (North-Holland, Amsterdam).

Diewert, W.E. and T.J. Wales, 1987, Flexible functional forms and global curvature conditions, Econometrica 55, 43–68.

Dilnot, A.W., J.A. Kay and C.N. Morris, 1984, The reform of social security (Oxford University Press).

Feenberg, D., 1982, The tax treatment of married couples under the 1981 tax law, NBER Working paper 872.

Hausman, J.A., 1979, The econometrics of labour supply with convex budget sets, Economics Letters 3, 171–177.

Hausman, J.A., 1980, The effects of wages, taxes and fixed costs in women's labour force participation, Journal of Public Economics 14, 161–194.

Hausman, J.A., 1981, Stochastic problems in the simulation of labour supply, NBER Working paper 788.

Hausman, J.A., 1985, The econometrics of nonlinear budget sets, Econometrica 53, 1225–1282.

HMSO, 1980, The taxation of husband and wife, Cmnd. 8093.

HMSO, 1986, Reform of personal taxation, Cmnd. 9756.

Hurd, M.D. and J. Pencavel, 1981, A utility based analysis of the wage subsidy program, Journal of Public Economics 15, 185–201.

Kay, J.A. and C. Sandler, 1982, The taxation of husband and wife: A view of the debate in the green paper, Fiscal Studies 3, 173–187.

King, M.A., 1983, Welfare analysis of tax reforms using household data, Journal of Public Economics 21, 183–214.

King, M.A., 1987, Empirical analysis of tax reform, in: T.F. Bewley, ed., Advances in econometrics (Cambridge University Press).

Meghir, C., 1985, The specification of labour supply models and the simulation of tax and benefit reforms, Unpublished Ph.D. thesis (University of Manchester).

Neary, J.P. and K.W.S. Roberts, 1980, The theory of household behaviour under rationing, European Economic Review 13, 25–42.

Pencavel, J., 1987, Labour supply of men: A survey, in: O. Ashenfelter and R. Layard, eds., Handbook of labour economics (North-Holland).

Stern, N.H., 1986, On the specification of labour supply functions, in: R.W. Blundell and I. Walker, eds., Unemployment, search and labour supply (Cambridge University Press).

Symons, E. and I. Walker, 1986, The reform of personal taxation: A brief analysis, Fiscal Studies, 38–48, May.