Fiscal Studies (1996) vol. 16, no. 3, p106-25

Tax Policy Reform: Why We Need Microeconomics

RICHARD BLUNDELL[1]

This is the text of the IFS Annual Lecture 1995 delivered at the Merchant Centre, London on 13 June 1995.

I. THE MOVE TOWARDS EMPIRICAL MICROECONOMICS

It is 20 years ago this July that the committee under the chairmanship of James Meade was set up by the Institute for Fiscal Studies to take a fundamental look at the UK tax structure. What it produced stands as a landmark in the history of tax policy analysis and was an important springboard for subsequent IFS research. The membership of the committee included three research 'secretaries',[2] two of whom were recent presenters of the IFS Annual Lecture — John Flemming and Mervyn King — and the other of whom was to become Director of IFS — John

[1] Institute for Fiscal Studies and University College London.
The author is grateful to James Banks, Andrew Dilnot, Alan Duncan, Christopher Giles, Paul Johnson, Costas Meghir and Graham Stark for comments and helpful encouragement in assembling the material for this lecture. It forms part of the research programme of the ESRC Centre for the Microeconomic Analysis of Fiscal Policy at IFS. Data from the Family Expenditure Survey made available by the Central Statistical Office through the ESRC Data Archive have been used by permission of the Controller of Her Majesty's Stationery Office. The usual disclaimer applies.
[2] Soon to become full members of the committee.

Kay. With such an active history in research in this area, it would seem reasonable to ask 'what's new?'.

Looking back, the 1970s can be seen as an era in which applied theory gained considerable influence on the design of tax reform. This is associated with the names of prominent British economists Atkinson, Mirrlees and, of course, Meade. The 1980s saw an important move toward empirical microeconomics. This reflected a dual stimulus: on the one side, from the important questions growing out of the applied theory research for which theory alone could not provide complete answers and, on the other side, from the dramatic increase in the accessibility of large micro-data sets. The subsequent move toward the careful use of cross-section survey data provided considerable insight into distributional questions and gave life to theories based on individual behavioural responses to tax policy reform. However, gaining a reliable picture of behavioural responses has been a more difficult nut to crack. Precisely why this has been the case and why it is so important to get behavioural responses right provide the motivation for this lecture.

One could argue that microeconomists and commentators in the media alike have been all too ready to pin their policy recommendations on simple cross-section correlations. However, mistaking causation for correlation can be dangerous. This is especially the case in fiscal policy research, where the targeting of policy often induces a correlation between actions and incidence that will prove misleading for policy 'counterfactuals'. How frequently do we see microeconomic policy conclusions — which by definition require some causal link — based on raw correlations?

Despite this, there must be something useful in the cross-section. After all, attempts to learn anything of real interest on tax policy reform from the time-series analysis of macroeconomic data have been, at best, illusive. What has been missing from microeconomic analysis is, none the less, what is routine good practice in macro analysis: that is, to track behavioural responses to past reforms. To do this, it is necessary to combine information on individuals as they pass through policy changes. Typically, observed responses to past reforms are clouded by the different macroeconomic environments in which the specific policy reforms took place. However, reforms do not usually affect every individual at the same time or in the same way. The aim of much new applied microeconomic policy research is to exploit this, in a sense drawing on the best of both time-series and cross-section analyses. This approach does not in itself necessarily require panel or longitudinal data. Most of the repeated cross-section data bases available in the UK and elsewhere will serve well for this purpose.

In this lecture, I take a look at three critically important tax and benefit reforms. The first concerns the debate over the degree to which new savings are generated by the introduction of tax-exempt special savings accounts in the UK. This example is used to point out the errors from drawing inferences about policy from simple cross-section correlation. A careful analysis of the

microeconomic evidence points to a large degree of shifting from other interest-bearing accounts on the introduction of these special savings accounts.

To really pin down the changes in individual behaviour that are attributable to policy reforms themselves, it is important to track policy reforms over time and to have certain groups that are effectively excluded from the reform. The reforms to work incentives in the UK over the 1980s are the second of the sets of reforms I consider and they provide a natural setting for this. In this lecture, they are used to display the dramatic change in estimated responses that occurs in moving from the cross-section correlation to the 'before and after' comparison. However, in an economic world absent of real controlled experiments, the before–after results turn out to be equally misleading. This is because they ignore important composition changes resulting from self-selection between the taxpaying and non-taxpaying groups. After all, the theory, if true, predicts that responses will occur! Individuals should self-select. Accounting for such composition changes is shown to produce believable estimates. The degree to which this matters for analysing responses to tax policy reforms is graphically illustrated using the recent debate on 'in-work' benefits.

Finally, the question is posed, 'is all the effort put into estimating responses worth it for welfare measurement?'. Using the 'heated' policy debate over the welfare consequences of extending the UK VAT base to include domestic fuel, the lecture concludes with a 'cost–benefit' appraisal of the benefits for welfare measurement of accurately measuring behavioural elasticities.

II. LEARNING FROM TAX REFORMS

Tax-Advantaged Savings: New Savings or Old?

Did the introduction of the tax-advantaged savings accounts in the late 1980s and early 1990s encourage a large degree of new saving? Such policy reforms were characteristic of fiscal reform in many countries during this period and this question is, not surprisingly, top of the agenda in many other countries. In the UK, it was the introduction of personal equity plans (PEPs) and tax-exempt special savings accounts (TESSAs) that has drawn most comment. Although their returns are less favourably treated than private pensions and housing, they none the less carry a tax rate of zero (see Banks, Blundell and Dilnot (1994) for a detailed analysis). So what of their impact on new saving? The argument here is not that a tax-advantaged return will fail to induce any new saving but rather 'how much of the recorded saving that exploits the tax advantage is simply a result of shifting from less-favoured assets?'. For a given level of tax subsidy, what is the best way to create new saving?

Needless to say, those individuals who took up TESSAs and PEPs typically also have higher overall savings. Table 1 confirms this, using data from the

TABLE 1

The Take-Up of TESSAs

Group	*Per cent*	*Group*	*Per cent*
No children	3.976	Children	2.032
Home-owner	5.244	Not home-owner	1.269
Aged under 30	1.034	Aged 45-64	3.406
Top income band	5.405	Other income band	3.175
Education beyond 18	5.383	No education beyond 16	2.825
Social class AB	6.126	Social class DE	1.320
Retired	5.080	Unemployed	1.412

Source: Banks, Blundell and Dilnot, 1994.

Financial Resources Survey:[3] individuals who can be expected to have a higher-than-average level of overall savings are also those most likely to have TESSAs. For example, among home-owners the take-up was over 5 per cent, whereas for non-owners it was a little over 1 per cent. Similar comparisons can be made for higher-educated, higher-income and retired individuals. However, this type of simple cross-section relationship cannot be used to suggest that TESSAs and PEPs had little impact on the level of other savings, therefore increasing overall savings.

Do recorded savings in TESSAs reflect new saving? As TESSAs are a close substitute for existing interest-bearing accounts (IBAs), would not we expect a large degree of shifting from other IBAs? Suspicion is cast on such inferences drawn from this simple correlation by the time series of micro evidence[4] presented in Banks, Blundell and Dilnot (1994), who report the high percentage of accounts (approximately 65 per cent) that are held at the maximum level — even from the date of their introduction.[5] This is surely a reflection of significant portfolio substitution across accounts by savers. The increased return over an interest-bearing account for a typical TESSA holder saving the maximum allowable would be about 1.5 percentage points per annum if held for the required five-year period. The high level of savings on the introduction of TESSAs is unlikely therefore to have reflected all new saving. The positive

[3] The FRS has been compiled privately by the National Opinion Polls since 1988. The characteristics of the data and the structure of asset holdings are described in detail in Banks, Dilnot and Low (1994).

[4] The data used in this analysis are the repeated cross-sections from the Financial Resources Survey. This is a quarterly survey covering the period 1988–92. TESSAs were introduced in the 1990 Finance Act and became effective from January 1991. They were available from most financial institutions that also provided standard interest-bearing savings accounts.

[5] There was a limit of £3,000 for savings in TESSAs in the first year and £1,800 for each subsequent year subject to an upper limit of £9,000 in total.

TABLE 2

Portfolio Substitution?
The Proportion of Wealth in TESSAs and other IBAs

Variable	*Proportions*	*With controls*
TESSAs	-0.8875 (31.48)	-0.8518 (31.94)
TESSAs available	-0.0243 (6.78)	-0.0186 (5.41)
National Savings	-0.8848 (156.99)	-0.9184 (163.36)
Wealth	—	-0.0001 (40.37)

Notes: t-ratios are given in parentheses. The dependent variable is the proportion of wealth in IBAs. TESSAs and National Savings also refer to the proportion out of total savings. The final column additionally includes a full range of socio-economic controls.
Source: Banks, Blundell and Dilnot, 1994.

correlation may simply reflect the fact that individuals with a strong taste for saving in TESSAs also have a strong taste for saving in other assets.

To ascertain the degree to which TESSAs act as a substitute for specific other IBAs, we can consider the relationship between proportions of wealth held in various assets. This is presented in Table 2. The proportion of wealth held in IBAs is tracked through time and across individuals as a function of a TESSAs-available dummy, the proportion of wealth held in TESSAs and the proportion of wealth held in National Savings accounts. Increases in the proportion of TESSAs lead to a fall in the proportion of other IBAs.

This analysis of proportions picks out the main substitute, but one would expect proportions between substitutes to move in opposite directions. What about controlling for the level of wealth? The final column of Table 2 controls for the total level of wealth (and many other socio-economic characteristics). This reduces the offsetting effect of TESSAs on IBAs, but not by very much. If IBAs were the principal source of existing savings from which funds were raised for the purchase of TESSAs, these results suggest no more than 15 per cent of TESSA saving represented ‘new’ saving.

III. WORK INCENTIVES: TRACKING BEHAVIOUR OVER PAST POLICY REFORMS

1. ‘In-Work’ Benefits

The recent debate, on both sides of the Atlantic, concerning the ability of `in-work’ benefits to help induce work effort by the low-paid provides a compelling

reason for the accurate measurement of labour supply responses (see Atkinson and Mogensen (1993) and Blundell (1992)). Without such measurement, the effects of these increasingly important components of the welfare system are virtually unknown. Family credit in the UK, soon to be piloted on workers without children, already contributes to the income of over half a million low-paid workers, with an exchequer cost rising to £1.6 billion. In the US, the earned income tax credit is set to be the largest call on the federal government's welfare expenditure to working- age families, outstripping AFDC (aid for families with dependent children) and food stamps, and rising to an overall cost of some $16 billion.[6] All aim to 'make work pay' for certain target groups of low-income benefit recipients. Will they succeed? The answer is not clear cut.

To give focus to this debate, consider the effects of the recent reform to family credit in the UK in which an additional incentive is given to those families with a worker working over 30 hours per week. A detailed characterisation of this reform and its impact on working behaviour is given in Duncan and Giles (1995). The likely impact on the budget constraint of a single-parent worker is given in Figure 1. This worker, drawn from the 1993 Family Expenditure Survey, currently works a little over 16 hours per week and receives family credit. The issue is whether the increased new benefit at 30 hours is sufficient to induce an increase in labour supply. Notice that under the pre-reform constraint, the 70 per cent withdrawal rate for family credit provides little inducement to work in excess of 16 hours. To analyse labour supply responses to the post-reform constraint, we need some simple microeconomics and a measure of wage and income elasticities for such individuals.

Individuals currently working *in excess* of 30 hours could be induced to reduce hours. What is more, for married women working part-time and with a partner[7] who is working in excess of 30 hours, the inducement is to reduce hours. The 30-hour rule acts as a pure income effect for them, as is illustrated in Figure 2. The issue is whether it is a sufficient inducement to make her quit employment.

This simple analysis of this reform points to a recurrent feature of 'in-work' benefits — there are offsetting effects on hours for workers which can easily dampen the expected impact of such reforms. The question is, 'how big are these effects?'. To answer this, we need an estimate of labour supply responses.

2. Cross-Section Correlation and Before–After Comparisons

The cross-section analysis of labour supply by married women and lone mothers typically shows a strongly positive wage elasticity. Recent evidence points to an upward bias in these estimates. Figure 3 shows the simple correlation between hours and gross wages for a sample of working married women with employed

[6] See Dickert, Houser and Scholz (1994).

[7] The family credit is paid direct to the mother by default.

FIGURE 1

Effect on Budget Constraint of the '30-Hour' Rule: Single Parent

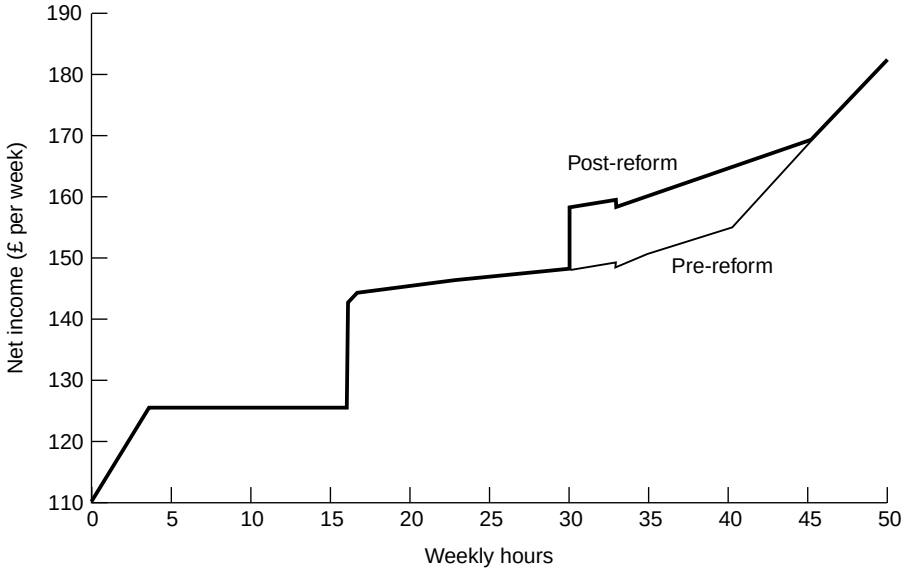

FIGURE 2

Effect on Budget Constraint of the '30-Hour' Rule: Married Woman

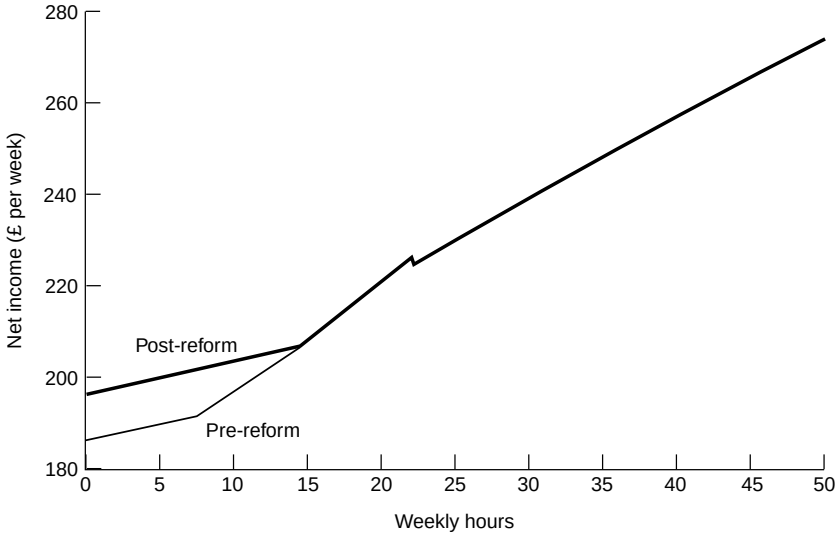

FIGURE 3

Labour Supply of Working Married Women with Employed Husbands

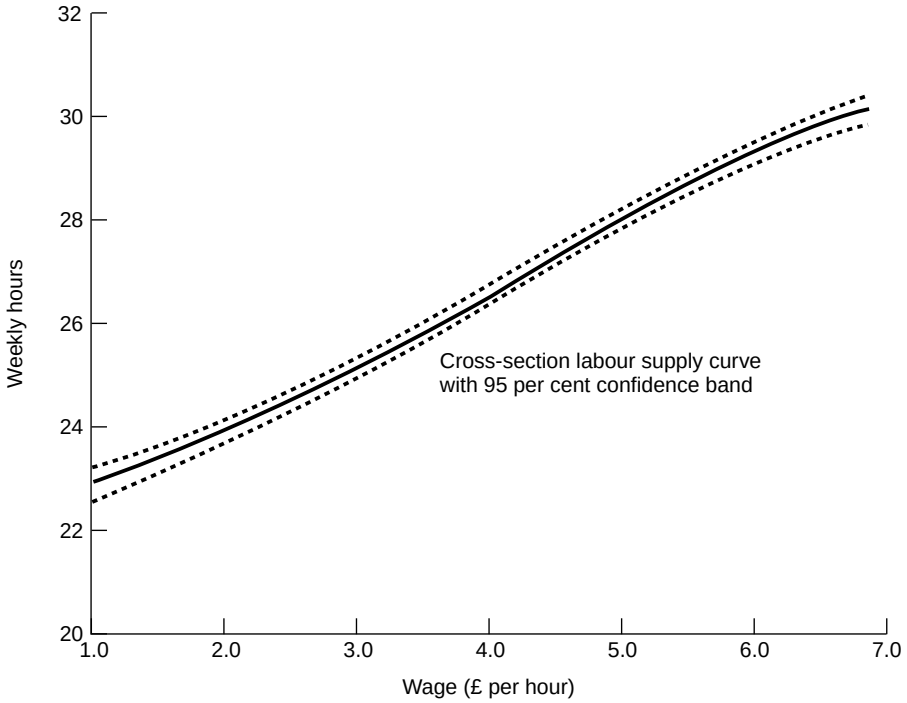

husbands from the FES.[8] Since these women have partners in employment, only a very small number are eligible for benefits. As a result, workers with higher earnings typically face higher marginal tax rates, and the use of gross hourly wages in the estimation of labour supply responses would lead to an upward bias in the response elasticity. However, the sensible alternative of using marginal wages requires a careful treatment of the budget constraint in cross-section analysis (see Burtless and Hausman (1978) and Hausman (1980), for example). Indeed, more recent evidence by MaCurdy, Green and Paarsch (1990) points to the potential for upward bias even after appropriate control for the tax system. Our approach at IFS[9] has been to relax the tight structure of the cross-section analysis, using the policy reforms over time to estimate more robustly the response effects.[10]

---

[8] A similar pattern can be found for working single parents — see Blundell, Duncan and Meghir (1992).

[9] See Blundell, Duncan and Meghir (1995).

[10] This is also designed to get round the endogeneity of gross wages as well as marginal wages. Mroz (1987) has shown the sensitivity of cross-section labour supply responses to different assumptions concerning the endogeneity of wages.

FIGURE 4

Differences in Hours and Wages between Taxpayers and Non-Taxpayers, 1978-92

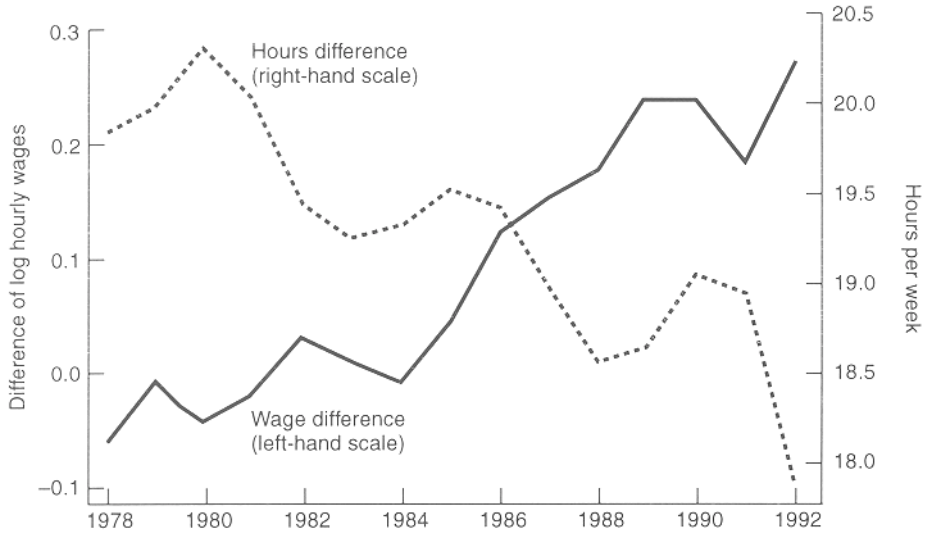

Given the tax and benefit policy reforms that have occurred in the UK during the 1980s, it might be that we could learn about labour supply responses by comparing individual behaviour *before and after* reforms. Indeed, since some individuals — those with earnings below the tax allowance — will be excluded from certain tax reforms, it might be expected that these could be used as a 'control' group to eliminate any general change in macroeconomic conditions or tastes that occurred. To examine this idea further, Figure 4 presents the time series of the difference in marginal wages between those paying taxes and those below the tax threshold. As could be expected, the graph shows a dramatic increase in this difference over time. So how did labour supply respond to this change in incentives? The time series of the corresponding hours differences is also plotted in Figure 4. Hours appear to have fallen relatively for those whose after-tax wages have increased, vividly turning around the simple cross-section result. But is this comparison believable? Not according to the results in Blundell, Duncan and Meghir (1995). They argue that a systematic change in the composition of the two groups would produce an important bias. Moreover, they argue that such changes occurred. Changes in composition occur for two reasons — self-selection as a result of the reforms (entrants and movers) and through tax allowance changes.

Adjusting for composition changes turns the difference of differences results around, producing a positive but smaller elasticity. These results are summarised in Table 3.

TABLE 3

Labour Supply Responses

Cross-section	*Difference of differences*	*Adjusting for composition changes*
0.41	0.49	0.15
(7.10)	(12.10)	(4.71)

Note: t-ratios are given in parentheses.
Source: Blundell, Duncan and Meghir, 1995.

3. Analysing the Reform to Family Credit

To examine the impact of behavioural responses to reforms to work incentives, we consider the likely reactions of individuals affected by the 30-hour rule reform to family credit described above. Figure 5 portrays the likely responses to the reform by the single parent analysed in Figure 1 above.[11] The pre- and post-reform indifference curves, highlighted in Figure 5, suggest that this woman would choose to increase her hours of work to 30. These indifference curves correspond to the higher 'cross-section' elasticity in Table 3. Figure 6 considers the same individual but with the lower elasticity. Notice that for this woman, there is much less chance of an increase in hours. Of course, for single parents working above 30 hours, the *higher* elasticity would predict more labour supply *reductions*.[12]

For a married woman with a working husband, the picture is very different. Figures 7 and 8 give an illustration. For the high-elasticity case, there is a larger incentive to move out of the labour market.

Of course, typical individuals are useful illustrations but are not likely to give an overall guide to labour supply responses. For this, we need to run a simulation across a representative sample of all families affected by the reforms. To do this, we use the IFS tax simulation model, SPAIN (which uses the IFS tax and benefit model, TAXBEN, to produce the budget constraint picture), using data from the 1993 Family Expenditure Survey. The summary results of this simulation are presented in Table 4. It is interesting to note the large number (10 per cent) of women in married couples affected by the reform who are predicted to reduce their hours. Most of these will leave the labour force. Among single parents, the majority who respond do increase their hours, but a large minority are seen to reduce their hours of work. These results are dampened when the more realistic lower-elasticity model is used, although the same overall features remain.

[11] Here we assume full take-up of family credit. This assumption can be relatively easily relaxed — see Dorset and Heady (1991), for example.

[12] This portrayal of responses by single parents to reforms to the UK family credit system is similar to that which followed earlier reforms to the hours rules – see Dilnot and Duncan (1992), for example.

FIGURE 5

Labour Supply Effect on Reforms to Work Incentives: Single Parent, 'High' Elasticity

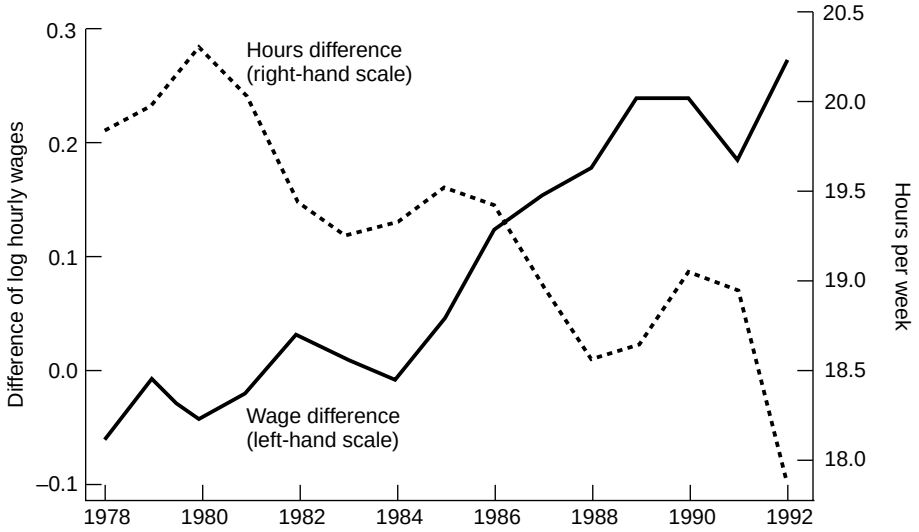

FIGURE 6

Labour Supply Effect on Reforms to Work Incentives: Single Parent, 'Low' Elasticity

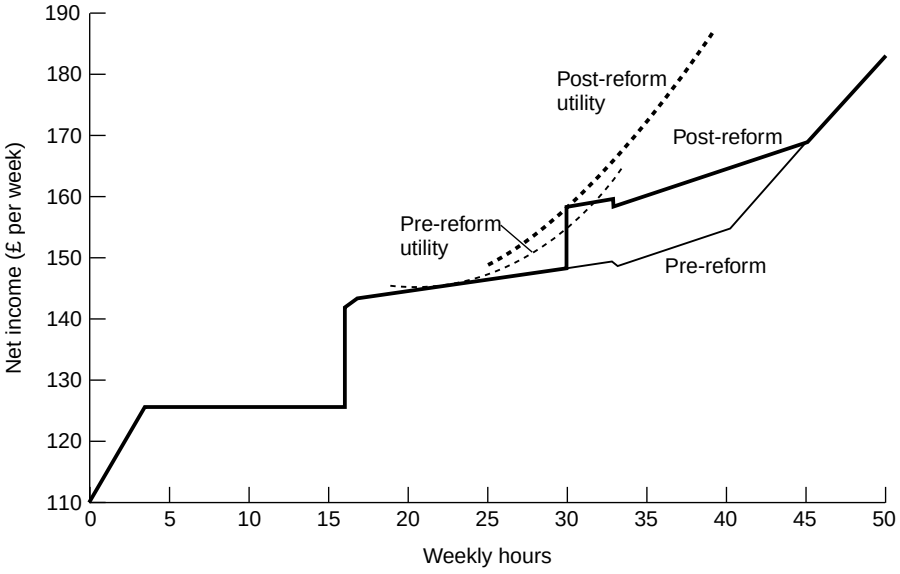

FIGURE 7

Labour Supply Effect on Reforms to Work Incentives: Married Woman, 'High' Elasticity

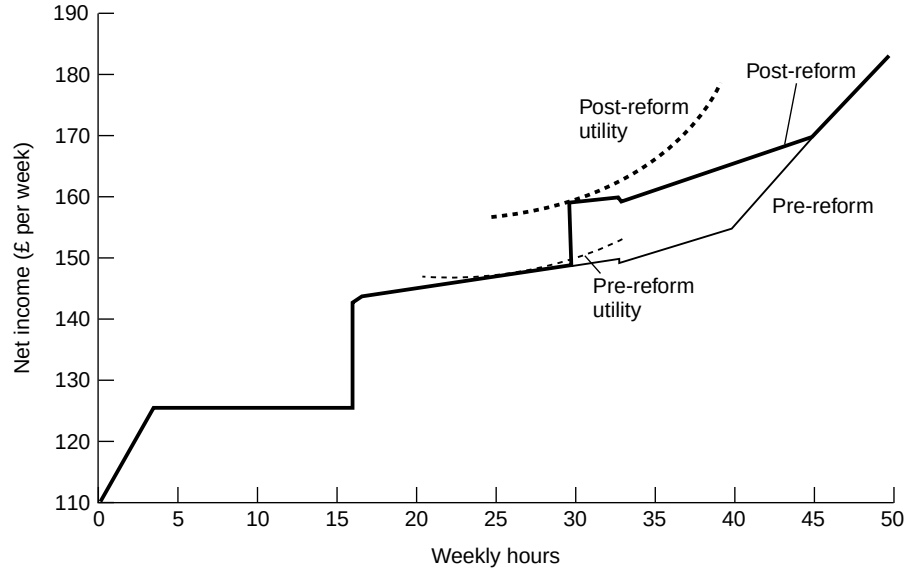

FIGURE 7

Labour Supply Effect on Reforms to Work Incentives: Married Woman, 'Low' Elasticity

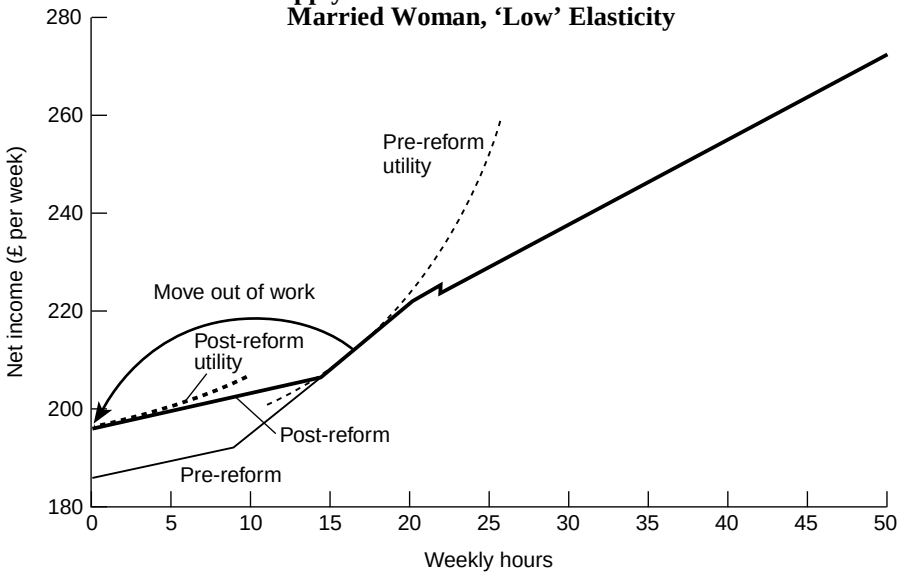

TABLE 4

The '30-Hour' Rule: Percentage Gains and Hours Changes

	Single parents		Couples	
	'High' elasticity	*'Low' elasticity*	*'High' elasticity*	*'Low' elasticity*
Percentage affected	£30%		30%	
Percentage gain	+5%	+4%	+1%	+3%
Hours change	-3%	-3%	-10%	-4%

Source: Duncan and Giles, 1995.

IV. MEASURING WELFARE GAINS AND LOSSES

Microeconomic policy research is as much about measuring welfare gains and losses as it is about predicting (or simulating) responses to reforms. These two objectives, however, are conceptually different. For example, in the labour supply context, a combination of income and substitution effects can produce relatively small responses to tax reform even though the change in dead-weight loss from taxation before and after the reform is significant. In this section, we seek to evaluate the gains and losses from using estimates of individual behavioural responses in the analysis of welfare gains and losses from tax policy reform. Clearly, if we want to predict behaviour, we need response elasticities. But, in terms of measuring the distribution of welfare gains and losses, some simple microeconomics of 'small' policy changes can avoid measuring elasticities (see Atkinson and Stiglitz (1980)). So what is the gain from engaging in the hard slog of estimating responses? To analyse this, we consider measuring the distribution of welfare costs of a reform that brings a new category of goods into VAT.

Evaluating welfare gains and losses requires information on individual preferences and individual utility weights in social welfare. Following Banks, Blundell and Lewbel (1994b), define the social welfare function

$$(1) \qquad \begin{aligned} U &= U(u_1,\ldots,u_H) \\ &= U(v_1(x_1,p),\ldots,v_H(x_H,p)) \end{aligned}$$

over households $h = 1,...,H$, where u_h is the attained utility level of household h, which equals the indirect utility function $v_h(x_h,p)$ of household h having total expenditure x_h and facing price p for the good or service under analysis. The indirect utility function also depends on attributes of the household such as demographic characteristics. Let $q_h = q_h(x_h,p)$ denote the quantity of the good purchased by household h, expressed as Marshallian demands; that is, as a function of prices and total expenditures.

As in Stern (1987, p. 54), for each household h we can use equation 1 to define a social marginal utility of income weight, θ_h. Then, for a small change in price, this change in utility can be approximated by

$$(2) \qquad \frac{\partial U}{\partial p} \approx -\sum_h \theta_h q_h .$$

This is commonly used to evaluate the social welfare effects of a price or tax change *without* explicitly estimating demand or indirect utility functions for individual households.

This 'first-order' approximation can be replaced by a more accurate second-order approximation,

$$(3) \qquad \frac{\Delta U}{\Delta p} \approx -\sum_h \theta_h q_h \left[\frac{\Delta p}{2} (\varepsilon_\theta + \varepsilon_q) + 1 \right]$$

where ε_θ is the price elasticity of the welfare weight and ε_q is the own price of demand. This shows that a second-order correction to the usual approximation depends on the price elasticities of both the Marshallian demands and the utility weights. Sensible welfare weights depend on prices except under very special circumstances (see Banks, Blundell and Lewbel (1994b) for precise cases), and therefore the price elasticity of θ_h in equation 3 is generally non-zero.

To avoid the issue of utility weights, consider the money metric measure of the change in social welfare. This is the amount of money required to get every household back to the same utility level they had before the price or tax change. The first-order approximation in this case is simply

$$(4) \qquad \frac{\partial X}{\partial p} \approx \sum_h q_h .$$

where X is the money metric measure of social welfare.

The second-order approximation simply includes the (compensated) price elasticity, ε_q^c:

$$(5) \qquad \frac{\Delta X}{\Delta p} \approx \sum_h q_h \left[\frac{\Delta_p}{2} \varepsilon_p^c + 1 \right]$$

FIGURE 9

Welfare Changes: 17.5 Per Cent VAT on Domestic Fuel

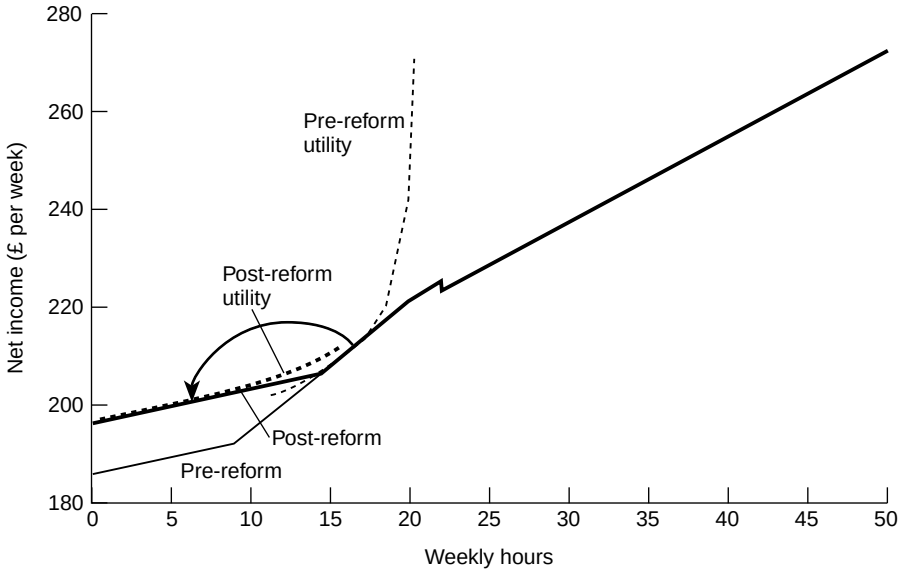

FIGURE 10

Percentage Welfare Changes: 17.5 Per Cent VAT on Domestic Fuel

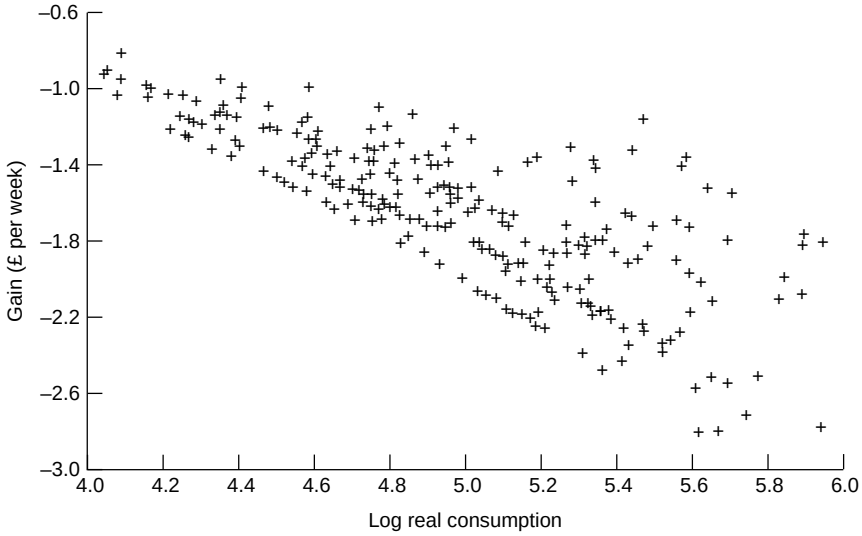

Apart from eliminating the terms in the utility weights, θ_h, the only difference between the second-order welfare approximation (equation 3) and the money metric approximation (equation 5) is that the compensated (Hicksian) own-price elasticity appears instead of the ordinary uncompensated (Marshallian) own-price elasticity.

Since compensated own-price elasticities are always negative, Banks, Blundell and Lewbel (1994b) note that we can unequivocally sign the bias in the first-order approximation of the money metric measure of price effects, in that the second- order approximation is always smaller in magnitude than the standard first-order approximation.

1. VAT on Fuel: The Distribution of Losses

Typically, the price or tax changes that are of the greatest policy interest are those involving *substantial* rather than marginal changes in price. In these cases, substitution effects can be non-trivial. To describe consumer behaviour, we will use the quadratic almost ideal model (QUAIDS). This simple extension of the Deaton and Muellbauer (1980) model is a budget share system that is quadratic in the logarithm of total expenditure, thus having the attractive property of allowing goods to have the characteristics of luxuries at low levels of total expenditure, say, and of necessities at higher levels.

This model is estimated in Banks, Blundell and Lewbel (1994a) using UK Family Expenditure Survey data for the period 1970–86. The budget system is defined over five goods — food, fuel, clothing, alcohol and 'other non-durable goods' — and the sample is restricted to non-retired married couples without children where the head is employed and the household lives in London or south-east England. These households are selected to form a reasonably homogeneous group in order to reduce the number of additional demographic factors that need to be controlled for in estimating preferences. Model parameters are estimated using the whole sample (4,785 observations over 68 quarterly price points) and elasticities are computed for each household. However, the welfare analysis that follows is carried out for only those observations observed in the final year of our data.

The tax change we choose to illustrate these approximations is a 17.5 per cent tax on domestic fuel. This represents a large price change, but is within the bounds of possibility in government tax reform. Indeed, 17.5 per cent is the current rate of value added (sales) tax in the UK, but many groups of goods are exempt; hence proposed moves towards a uniform expenditure tax would require tax changes of this magnitude. Figure 9 shows the money metric measure of welfare loss to individuals in our sample of the FES.

For this illustration, we have taken a sample of couples without children in which the head is in work. This is likely to be a relatively rich group. Figure 10 calculates the distribution of losses as a percentage of the individual household's

total consumption level.[13] For the poorest in this group of households, losses reach around 2 per cent of total consumption, confirming the results of Crawford, Smith and Webb (1993).

2. The Value of Measuring Elasticities

There is an obvious attraction to simply using information on observed commodity demands to assess the welfare implications of tax reform. No response parameters are required and therefore the analysis can transcend misspecification of preferences and is not subject to estimation error in own- or cross-price demand elasticities. However, tax reforms are often far from marginal and can involve a significant realignment of relative prices. In such cases, it might be argued that using information on elasticities should produce improvements in welfare measurement.

Figure 11 shows that the magnitude of the first-order approximation error in $\Delta X/\Delta p$ is considerably greater than that of the second-order error. On average, the first-order error is around 8 per cent, whereas the addition of the second-order term in the approximation reduces this average error to 0.1 per cent of the true welfare change.

These results are more extensively studied in Banks, Blundell and Lewbel (1994b), where it is shown that the difference between first- and second-order approximations has the same sign for every individual and therefore will not average out in any standard aggregate social welfare measure. For a tax reform that adds a new group of goods to the tax base in the UK at a tax rate of 10 per cent, we find that the bias can be of the order of 5 per cent. For smaller reforms, we show that suitable first-order approximations can work very well. However, the second- order approximations in our examples uniformly produce improvements in the measurement of changes in aggregate social welfare.

V. PUTTING IT ALL TOGETHER

The aim of this lecture has been to show how we can best exploit micro data and microeconomics for tax policy analysis. We have argued that microeconomists and commentators in the media alike have been too ready to pin their policy recommendations on simple cross-section correlations. The spurious causation that results can prove highly misleading. Using three recent UK tax reforms, the lecture has highlighted three important aspects of empirical microeconomics.

The debate over the degree to which new savings are generated by tax-exempt savings accounts was used to point out the errors from simple cross-section correlation. Although the simple correlation between savings in TESSAs and total saving is strongly positive, suggesting little evidence of a reduction in

[13] For arguments as to why consumption is a good measure of welfare, see Blundell and Preston (1994).

FIGURE 11

Approximation Errors as a Percentage of True Changes

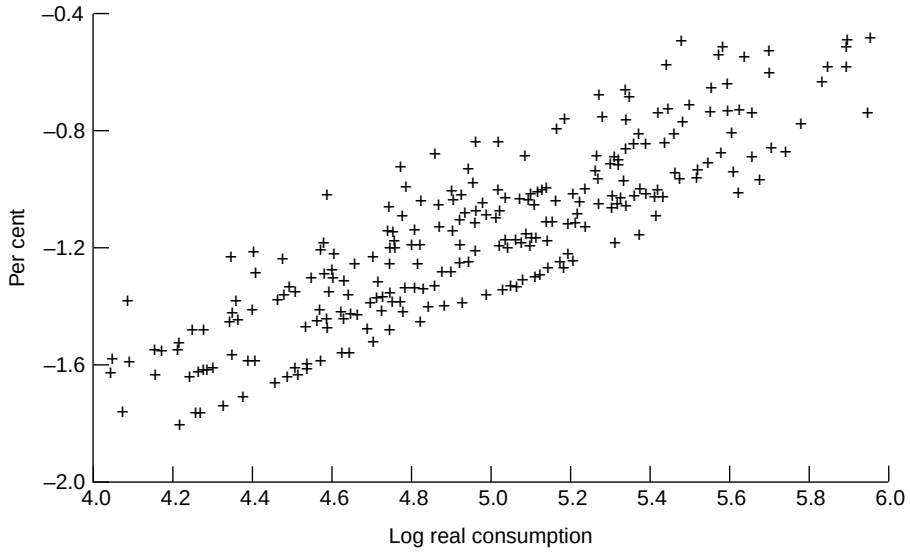

other saving on the introduction and expansion of TESSAs, a careful analysis of the microeconomic evidence pointed to a large degree of shifting from other interest- bearing accounts. It seems that no more than 15 per cent of total TESSA savings could be attributed to new savings.

To really pin down the impact of policy reforms, however, it is important to track responses to a large system sequence of policy reforms. The reforms to work incentives over the 1980s provided a natural setting for this. In this lecture, they were used to display the dramatic change in results that can occur in moving from the cross-section correlation to the analysis of the time series of policy reforms. The degree to which this matters was graphically illustrated using the recent debate of the effectiveness of `in-work' benefits to encourage work and help the low-paid.

It is not just behavioural responses that are of interest; the impact of reforms on individual welfare and its distribution is of considerable importance. It turns out that for 'small' reforms, the welfare effects of a tax reform can be analysed essentially ignoring behavioural effects. However, tax reforms are typically far from 'small'. Using the debate over the distributional consequences of extending the UK VAT base to include domestic fuel, the lecture concluded with a 'cost–benefit' appraisal of the benefits of accurately measuring responses. For reforms of this type, it was shown that knowledge of behavioural responses considerably improves our ability to measure the welfare gains (or losses) of the reform.

REFERENCES

Atkinson, A. and Mogensen, G. (eds) (1993), *Welfare and Work Incentives: A North European Perspective*, Oxford: Clarendon Press.

— and Stiglitz, J. (1980), *Lectures on Public Economics*, McGraw-Hill.

Banks, J., Blundell, R. and Dilnot, A. (1994), 'Tax-based savings initiatives in the UK', paper presented at OECD conference on international comparisons of household saving, Paris, May.

—, — and Lewbel, A. (1994a), 'Quadratic Engel curves, indirect tax reform and welfare measurement', Univerity College London, Discussion Paper no. 94-04.

—, — and — (1994b), 'Tax policy reform: do we need demand system estimation?', Institute for Fiscal Studies, Working Paper no. 94/12.

—, Dilnot, A. and Low, H. (1994), *The Distribution of Wealth in the UK*, Commentary no. 45, London: Institute for Fiscal Studies.

Blundell, R. (1992), 'Labour supply and taxation: a survey', *Fiscal Studies*, vol. 13, no. 3, pp. 15–40.

—, Duncan, A. and Meghir, C. (1992), 'Taxation and empirical labour supply models: lone parents in the UK', *Economic Journal*, vol. 102, pp. 265–78.

—, — and — (1995), 'Estimating labour supply responses using tax policy reforms', Institute for Fiscal Studies, Working Paper no. 95/9.

— and Preston, I. (1994), 'Income or consumption in the measurement of inequality and poverty?', Institute for Fiscal Studies, Working Paper no. 94/12.

Burtless, G. and Hausman, J. (1978), 'The effect of taxes on labour supply', *Journal of Political Economy*, vol. 86, pp. 1103–30.

Crawford, I., Smith, S. and Webb, S. (1993), *VAT on Domestic Energy*, Commentary no. 39, London: Institute for Fiscal Studies.

Deaton, A. and Muellbauer, J. (1980), 'An almost ideal demand system', *American Economic Review*, vol. 70, pp. 312–36.

Dickert, S., Houser, S. and Scholz, J. (1994), 'The earned income tax credit and transfer programs: a study of labor market and program participation', in J. Poterba (ed.), *Tax Policy and the Economy*, vol. 9, Cambridge: MIT Press, forthcoming.

Dilnot, A. and Duncan, A. (1992), 'Lone mothers, family credit and paid work', *Fiscal Studies*, vol. 13, no. 1, pp. 1–21.

Dorset, R. and Heady, C. (1991), 'The take-up of means-tested benefits by working families with children', *Fiscal Studies*, vol. 12, no. 4, pp. 22–32.

Duncan, A. and Giles, C. (1995), 'Benefit reforms and labour supply incentives in the UK: the family credit', paper presented to the Royal Economic Society Special Policy Session, University of Kent, March; mimeo, Institute for Fiscal Studies.

Hausman, J. (1980), `The effect of wages, taxes and fixed costs on women's labor force participation', *Journal of Public Economics*, vol. 14, pp. 161–94.

MaCurdy, T., Green, D. and Paarsch, H. (1990), 'Assessing empirical approaches for analysing taxes and labour supply', *Journal of Human Resources*, vol. 25, pp. 415–90.

Mroz, T. (1987), 'The sensitivity of an empirical model of married women's hours of work to economic and statistical assumptions', *Econometrica*, vol. 55, pp. 765–800.

Stern, N. (1987), 'The theory of optimal commodity and income taxation', in D. Newbury and N. Stern (eds), *The Theory of Taxation for Developing Countries*, Oxford: Oxford University Press.