

TAX CUTS, TAX REFORM AND LABOUR SUPPLY RESPONSES

TIM CALLAN, ARTHUR VAN SOEST, JOHN WALSH*

1.
Introduction

How strong are the links between income tax cuts and employment growth? Does the employment impact of a tax cut depend on how income taxes are cut? These are central questions for income tax policy. One critical factor in each case is the response of the supply of labour to income tax cuts: for a given demand curve,¹ the more responsive (elastic) is labour supply, the greater the ultimate impact on employment. An adequate examination of labour supply requires detailed analysis of the choices made by a representative sample of real individuals. Studies of this type, exploring the different dimensions of labour supply – decisions regarding participation and non-participation as well as hours of work – have become widespread in the international literature,² but labour supply analyses which can take account of the key features of the tax/transfer system are, as yet, sparse in the Irish context.

In this paper we summarise the key findings of a recent study (Callan *et al.*, 2003) and draw out the implications for the links between tax cuts and employment growth. The analysis focuses on married men and women: the former are typically among the least responsive groups and the latter among the most responsive. Other responsive groups include migrants and potential migrants (including many young, single people): these fall outside the scope of the present work, and a rather different framework might be required to model their opportunities and choices.³ The analytical framework involves modelling each individual's budget constraint –

* Tim Callan is a Research Professor and John Walsh is Research Analyst with The Economic and Social Research Institute. Arthur van Soest is a Senior Economist at the RAND Corporation.

¹ Assuming that labour demand is not perfectly inelastic.

² See, for example, the review by Blundell and MaCurdy (1999).

³ Lower tax rates could also lead to reductions in tax evasion, including reallocation of labour supply from illegal untaxed activity into taxed activity. Again, this is outside the scope of the present paper.

how much disposable income would they end up with for various choices of hours of work. We know their preferred choice, given the set of options available at the time of interview. From this it is possible to infer an underlying preference (or “utility”) function which can depend on a range of personal and family characteristics. Tax changes alter the choices available, but the preference function can be used to identify the new optimum that the individual will choose.

The structure of the paper is as follows. Section 2 sets out the basic framework and results on the responsiveness of married men and married women to changes in their gross wages. Section 3 briefly describes the framework for exploring the impact of a tax change on labour supply. Section 4 examines the impacts of alternative forms of tax cut. Section 5 considers how the labour supply decisions of married men and women are affected by changes in the income tax treatment of couples.

2. A Framework for the Analysis

How can we derive estimates of labour supply responsiveness which can be used to analyse the impact of changes in tax and benefit policy? We need to be able to separate the impact of individual preferences, on the one hand, and budget constraints, on the other, in determining hours of work (see Blundell, 2001). Policy reforms which change the budget constraint facing individuals and families can then be simulated. This separation of the influence of budget constraints (influenced by tax and welfare policies) and individual preferences can be achieved using what is known as a discrete choice structural model of labour supply. The “discrete choice” element means that we estimate labour supply in blocks of, say, 4 or 8 hours per week, rather than at a very fine level of detail. This is particularly important when modelling the behaviour of couples, where estimation must search over a grid of all possible hours combinations for husband and wife. By focusing on the modelling of labour supply to the nearest 4 hours, say, the search grid can be made more manageable, which allows greater attention to be paid to the complexity of the budget constraint.

A full description of the methods used can be found in Callan *et al.* (2003). Here we outline the broad rationale behind the approach, and present some of the labour supply elasticities calculated using the estimates.

Individuals are observed as having chosen certain hours of work, and the hourly gross wage that they actually earn can be calculated if they are in employment, or estimated if they are not presently in paid employment. The net family incomes they would obtain at other hours of work can then be simulated using routines for the calculation of taxes and benefits adapted from a tax-benefit model (in this case, *SWITCH*). The fact that each individual has chosen the observed package of income and hours of work is taken as indicating that this is the one which is preferred to all the other

feasible choices, and maximises the family's overall welfare.⁴ Information of this type on a large sample of families can be used to identify the underlying influences on family welfare which guide labour supply choices. Once these have been identified, it becomes possible to predict how individuals may behave if tax or welfare policy is changed in certain ways – as will be seen in Sections 3 and 4.

Labour Supply Estimates

The discrete choice structural labour supply model constructed by Callan *et al.* (2003) captures important features of household labour supply behaviour from a policy point of view. The benchmark model approximates the budget constraint using multiples of 8 hours i.e., 0, 8, 16, 24, 32, 40 and 48 hours. This yields a grid of (7 x 7 =) 49 possible combinations for a couple. Previous work by van Soest (1995) and Callan and van Soest (1996) found that estimates of labour supply elasticities were not much affected if a finer grid was used instead.

The model accounts for the full structure of the basic tax rules and the key feature that the social welfare system provides a floor to income. It captures simultaneously the participation decision and the decision on hours worked, by allowing for fixed costs of work. It takes account of the fact that wage rate information is not directly available for those who are not in employment.

The model was estimated using data from the 1994 Living in Ireland Survey, which included questions on desired hours of work from which it was possible to construct a measure of “preferred hours” for those in work, unemployed or out of the labour market. (For details, see Callan *et al.*, 2003). This information, collected only in the 1994 wave of the Living in Ireland Survey, is particularly valuable in trying to identify individuals' preferences as regards hours of work. The key focus of the study is on identifying the influences of preferences, on the one hand, and opportunities, shaped by labour market forces and tax/transfer policy, on the other. Once this is done, it is possible to simulate the impact of various policy changes on labour supply decisions.

Estimates of elasticities of labour supply for married men and married women are shown in Table 1. In this table, the (own- or cross-) wage elasticity of labour supply of a given group of people (husbands or wives) is defined as the percentage change in total desired hours of that group if all the before-tax wage rates (of husbands or wives) in that group rise by 1 per cent. Note that all gross wage rates change by 1 per cent, but the tax system remains unchanged. The way in which net wage rates change is thus not fixed *a priori*, but driven by the existing tax system. On average,

⁴ It is, of course, possible that actual hours worked diverge from the individual's preferred hours at the given hourly wage. The estimates discussed here use information specially gathered on individuals' preferred hours of work.

after-tax wage rates will change by slightly less than 1 per cent, due to the progressive nature of the tax rules.

Many studies only consider or report elasticities for the average (“representative”) family. In a highly nonlinear model, however, these elasticities may not be a good guide to the consequences of wage changes for a heterogeneous population. Other studies consider average elasticities instead of elasticities of the average. The average elasticity can be seen as a weighted aggregate elasticity of hours worked, where more weight is given to people with lower desired hours. Some studies look at elasticities of hours worked *conditional upon participation*. For policy analysis, however, the effect on participation is at least as important as the effect on hours worked given participation, particularly for married women. The elasticities reported here take full account of the (positive) impact of the wage rate on the participation decision (with desired hours equal to zero for non-participants).⁵

Table 1: Labour Supply Elasticities for Married Men and Married Women with Respect to Wage Changes

Change in:	Elasticity of average preferred hours to change in wages	
	Husbands	Wives
Male wage	0.25	-0.35
Female wage	-0.07	0.88
Both wages	0.18	0.48
<i>Note:</i> A 1 per cent rise in the male wage leads to a 0.25 per cent rise in average preferred hours of married men, and a fall of 0.35 per cent in the average preferred hours of married women		
<i>Memo item:</i> Average preferred hours for men was 35.8 and for women 11.1 hours per week.		

For men, the own-wage elasticity is estimated at 0.25. That is, if all gross wage rates of the men in the sample increased by 1 per cent, with women’s wage rates remaining unchanged, the total desired hours of all men would increase by 0.25 per cent. Most of this effect is due to increased participation: a rise in each husband’s gross wage rate of 1 per cent would induce an increase of the number of men willing to participate of almost 0.2 percentage points, i.e., by 0.21 per cent of the actual participation rate of almost 90 per cent. For women, the estimated own-wage elasticity is 0.88. The elasticity of the participation rate is 0.49, which again explains the largest part of the total labour supply elasticity. These estimates are well in line with the broad range of empirical findings of labour supply elasticities for other countries (see, for example, Killingsworth and Heckman, 1986), even though, as explained above, a comparison is hampered by the fact that the large number of empirical studies are based on an almost as large number of elasticity concepts. When estimated cross-wage elasticities (-0.07 for men, -0.35 for women) are taken into account, the model predicts

⁵ Like the vast majority of labour supply elasticities quoted in the literature, these do not take account of wider repercussions in the labour market of the changes in labour supply.

that a general increase in wages of 1 per cent would see desired hours rise by 0.18 per cent for men, and by 0.48 per cent for women.

Thus, in line with other findings for Ireland (e.g., Doris, 2001) and internationally, the labour supply of married women is significantly more responsive to an increase in their wage rate than men (with respect to the male wage rate). For both men and women, increased participation accounts for the major part of the response, with increases in hours of work playing a lesser role.

The analysis deals with desired or preferred hours of work at the wage rate the individual currently commands. This allows for considerable simplification over analyses which must deal with the potential for involuntary unemployment or actual hours of work which diverge from preferred hours. It can also be seen as allowing for maximum flexibility in labour market response. In some circumstances changes in desired hours will not translate into changes in actual hours because of constraints on individual behaviour (e.g., having to choose between full-time and part-time work, or being involuntarily unemployed). Nevertheless, it is of interest that the own-wage elasticities for men and women reported by Callan and van Soest (1996), based on 1987 data on actual hours and incorporating modelling of involuntary unemployment and constraints on hours, are quite similar to those reported here (0.15 for men, 0.67 for women).

3. Labour Supply Response to Tax Policy Changes

Callan *et al.* (2003) analyse the impact of changes in tax rates, tax free allowances and wider rate bands on labour supply, as well as the impact of increased independence in the taxation of husbands and wives. The structural model is particularly useful for this purpose, since it accounts for the complete structure of the tax system and cash benefits, including any “kinks” or “jumps” in budget constraints.⁶ Moreover, the model predicts the effects on labour market participation – whether or not an individual seeks paid work – as well as the distribution of hours worked for those in paid employment.

The way in which the effects are predicted is very similar to the method of computing the elasticities in Section 2. Using the parameter estimates, the labour supply of each individual under the actual 1994 tax rules is predicted. The simulation is then repeated using the tax rules after the reform in question. Comparing the two outcomes gives the predicted changes. For the simulation after the reform, it is assumed that gross (pre-tax) wage rates remain the same. Thus, general equilibrium effects are not taken into account.

In Section 4 we compare the labour supply effects of equal-valued tax cuts through the four main routes actually used to cut

⁶ The analysis, like that in most international studies, does not include non-cash benefits such as the medical card. For a path-breaking US study analysing participation in both cash and non-cash programmes see Keane and Moffitt (1998).

taxes over the past decade: the top and standard rates of tax, the width of the standard rate tax band, and the size of the basic personal allowance. Section 5 considers the impact of a structural reform, in which the transferability of allowances between husbands and wives is reduced. Alternative uses of the incipient rise in revenue – increasing child benefit, or reducing income tax rates – are contrasted.

4. Analysis of Labour Supply Response to Alternative Tax Cuts

4.1 SPECIFICATION OF ALTERNATIVE TAX CUTS

In this section, we ask which form of tax cut does most to stimulate aggregate labour supply? In order to answer this question, we simulate the labour supply response to four different types of income tax cut: a cut in the standard rate of tax, a cut in the top rate of tax, a rise in the personal allowance, and a widening of the standard rate band. Each form of tax cut is scaled to have approximately the same Exchequer cost on a static basis i.e., before any behavioural response. Thus, the options considered are:

1. A cut of 2.8 percentage points in the *standard rate of tax* (from 27 per cent to 24.2 per cent).
2. A cut of 6.3 percentage points in the *top rate of tax* (from 48 per cent to 41.7 per cent).
3. An increase of IR£2,400⁷ in the *standard rate band* (from IR£8,200 to IR£10,600 for single taxpayers, and double these limits for married couples).
4. An increase in the *basic personal allowance* (and the allowance for widowed/lone parents) of IR£465⁸ (from IR£2,350 to IR£2,815).

Our exploration of the labour supply impact of these different forms of tax cut is of considerable interest in its own right. In addition, however, it helps to shed light on the wider debate concerning the appropriate balance between these different elements when cutting income taxes. Between 1987 and 2002 considerable resources were devoted to a range of tax cuts, over and above what would have been needed for indexation of the income tax system with respect to prices or wages. Over this 15 year period the standard rate of tax was cut from 35 per cent to 20 per cent, and the top rate of tax was reduced from 65 per cent to 42 per cent. While wages almost doubled over the period, the basic personal allowance was almost trebled and the standard rate band for a single person was increased by almost 470 per cent.

There was considerable debate as to the appropriate structure of tax cuts at various times during this period. While such debate usually received greatest attention at annual budget time, it was also a recurring theme in partnership negotiations, and was a leading issue in the 1997 election campaign. On the one hand, it was argued that a focus on increasing personal allowances would help to

⁷ In euro terms this comes to approximately €3,050.

⁸ In euro terms this comes to approximately €590.

concentrate the benefits of tax cuts on low-income earners. As against this, it was sometimes argued that other forms of tax cut – including widening of the standard rate band and reduction of income tax rates – could lead to greater employment growth, and might therefore be preferable. The framework outlined here allows these issues to be analysed in a more comprehensive way than heretofore.

4.2 ESTIMATED LABOUR SUPPLY RESPONSES TO ALTERNATIVE TAX CUTS

Table 2 shows the estimated changes in participation rates for men and women in response to various tax policy changes. A cut in the standard rate of tax of just under 3 percentage points leads to a rise of about half a percentage point in both male and female participation rates. An increase in personal allowances (with similar Exchequer cost) has very similar effects. Changes to the top rate of tax or to the standard rate band, however, have rather different consequences. A cut in the top rate of tax of more than 6 percentage points is estimated as leading to a rise of about 1 percentage point in the participation rate for married women, but only to a very small rise in the participation rate for married men (0.1 percentage points). Widening the standard rate band leads to similar, but marginally greater changes in participation rates.

Table 2: Response of the Labour Force Participation Rate to Selected Tax Cuts

Tax Cut Option	Change in Husbands' Participation Rate	Change in Wives' Participation Rate
Standard rate cut by 2.8 percentage points	+0.5	+0.6
Top rate cut by 6.3 percentage points	+0.1	+1.0
Standard rate band up by Ir£2400 from Ir£8200	+0.2	+1.1
Personal allowances increased by Ir£465 from Ir£2,350	+0.5	+0.5

Note. All options had an Exchequer cost of about IR£200m per annum (€250m), including the impact on single persons.

The overall change in participation (male and female combined) is rather similar across the different options. The major difference is in the sex distribution of the change in participation. A cut in the top rate of tax, or a widening of the standard rate band, prompts a greater increase in female participation and much less in male participation. A cut in the standard rate of tax, or an increase in personal allowances, leads to similar increases in participation rates for both sexes. There may, of course, be other differences in the impact of policy (e.g., as between high and low income families, or those with above and below average education) which could be of interest.

The overall labour supply response includes not only changes in labour market participation, but also changes in desired hours for those who were initially in employment or seeking employment. As

with the participation response, the aggregate (male and female combined) labour supply response is rather similar across the options, with a rise in average desired hours of 0.2 to 0.3. The differences between the options are also similar to those observed for the participation response: a cut in the top rate of tax or a widening of the standard rate band give rise to a much greater response by married women than by married men. There is little difference between the sexes in the response to a standard rate tax cut or a rise in personal allowances.

5. Tax Treatment of Married Couples

The second major area examined here relates to the income tax treatment of couples. Over time a number of countries have moved from systems involving “income-splitting” or extensive transferability of allowances between spouses to systems involving greater independence in the tax treatment of husband and wives – and, correspondingly, more restricted transferability of allowances and/or bands.⁹ More recently the Irish tax system has moved towards greater independence in the tax treatment of couples, in what has been termed “individualisation” of the standard rate tax band. There has been considerable speculation about the likely impact of this change on the participation of married women in the paid labour market. Analysis of the type set out here is necessary to provide estimates of likely impacts which can be used to inform the debate.

The Irish tax system – like the UK system – initially treated married couples as a unit for income tax purposes, with the wife’s income being aggregated along with that of her husband. While there was a “married man’s allowance” tax was assessed on the basis of the same band width as for single persons. Compared to two cohabiting single persons, a married couple received a marriage subsidy if the wife was not earning an independent income, or earned a very low one. But if the wife’s earnings were greater, she, and the couple, faced a substantial tax penalty – a married couple with both partners in employment could face a higher tax bill than an unmarried couple in identical circumstances.

The Supreme Court ruled in the Murphy case (1979) that this feature of the tax system was unconstitutional. A number of responses to this ruling may have been possible. The one chosen by the government, and implemented in Budget 1980, was to allow doubled rate bands and doubled allowances to all married couples. Formally, this was equivalent to allowing “income splitting” i.e., calculating the couple’s tax liability on the basis of assigning half the income to each partner and taxing them as if they were single. It was also equivalent to full transferability not only of allowances but also of rate bands. Married couples were permitted to minimise their tax

⁹ See OECD, 1977; Callan *et al.* (2001). O’Donoghue and Sutherland (1999) found that 10 out of 15 EU countries had income tax systems which were based around independent or individual taxation of husbands and wives.

liabilities by assigning allowances and rate bands freely to either partner.

The main reason given for this approach in the 1980 Budget Speech was that:

A narrow approach towards effecting the minimum changes to meet the Supreme Court's decision would lead to unjustifiable discrimination against the one-income family, particularly where a married woman elects to care for the family on a full-time basis at home rather than take up work outside the home. (Minister for Finance, 1980, p. 18).

Callan and Farrell (1991) comment that if, as would appear from this statement, the policy objective was to subsidise childcare undertaken by married women in the home, the mechanism chosen was a rather inefficient one. The tax subsidy is not conditional on having children, but on being married – implying that “... the benefit from this tax break is, in terms of its main stated objective, rather inefficiently targeted”. Furthermore, the mechanism imposed high effective tax rates on married women with and without children, thereby giving rise to a substantial efficiency loss. Callan and Farrell concluded that other methods of providing child income support, notably through child benefit, might involve smaller efficiency losses and better targeting. Fahey (1998) came to a similar conclusion, based on an analysis of Labour Force Survey data, finding that “Many who receive the subvention are not engaged in childcare, and many of those with young children who have a heavy childcare burden do not receive the subvention”.

The major structural innovation in Budget 2000 was a move towards individualisation of the standard rate tax band. This involved restricting the extent to which tax bands are transferable between spouses. In 1999 the standard rate band was IR£14,000 for an individual, or IR£28,000 for a couple i.e., a non-earning partner could transfer 100 per cent of his or her tax band (and, indeed, of his/her allowance). In 2000, full transferability of tax allowances remained as before, but there were, in effect, restrictions on the transferability of the standard rate band. The band for a single person was increased from IR£14,000 to IR£17,000 per annum; for a married couple with one income the band remained unchanged at IR£28,000 per annum; but the band for a married couple, both earning, rose to IR£34,000 (twice the single band, thereby meeting the requirement of “no marriage penalty”). Thus, in effect, only two thirds $[(IR£28,000 - IR£17,000) / 17,000 = 11/17]$ of a non-earning partner's band was transferable.¹⁰ The stated objective was to arrive at a position after three years where each individual, whether single or married, has his/her own standard rate tax band which can be set off against his/her own income but cannot be transferred between spouses. By December 2001 the proportion of

¹⁰ In the immediate aftermath of the budget, a special Home Carer's Allowance was introduced for couples with one partner staying at home to care for a child or children, an elderly person or someone with a disability.

the band which was transferable had fallen to about one-third, remaining at that level after Budget 2003.

In this section, we do not attempt to summarise the extensive, and often rather heated, debate that has grown up around the shift from an income-splitting system to a system with greater independence (and less transferability of rate bands) between husband and wife. Our aim instead is to bring new evidence on the likely outcomes linked with different policy choices in this area, which can help to inform those on all sides of the debate. With this in mind, we examine the potential size of labour supply responses to a full-scale individualisation of tax bands, and how this is affected by alternative uses of the incipient rise in tax revenue associated with restrictions on the transferability of the rate band.

5.1 LABOUR SUPPLY RESPONSE TO ALTERNATIVE REFORM PACKAGES

Table 3 shows the impact of alternative ways of implementing increased independence in the tax treatment of husbands and wives. Option (A) simply involves the elimination of transferability of the standard rate tax band, and would generate something over £200m per annum in extra tax revenue.¹¹ Option (B) returns this revenue to taxpayers, via proportionate cuts in the standard and top tax rates. Option (C) is also revenue neutral, but the incipient rise in revenue is used to fund an increased child benefit.

Table 3: Response of Husbands' and Wives' Participation Rates to Increased Independence in Tax Treatment of Married Couples

Change in Tax Structure	% Point Change in Husbands' Participation Rate	% Point Change in Wives' Participation Rate	Net Change in Exchequer Revenue as Estimated by SWITCH on Full Sample ¹²
(A) Standard rate band made non-transferable	-0.5	+1.8	+Ir£210m
(B) Band non-transferable, tax rates cut to 25.4% and 45.1%	-0.1	+2.6	-Ir£8m
(C) Band non-transferable, Child Benefit increased by 69%	-0.9	+1.6	+Ir£1m

A notable feature of option (A) is that it gives rise to a net increase in labour market participation (a fall in married men's participation being more than offset by a rise in the participation of married women), while at the same time actually increasing net revenue for the Exchequer. Options (B) and (C), returning this revenue via general tax cuts or via child benefit, are designed to be

¹¹ All calculations are undertaken in a 1994 setting.

¹² In euro terms, the exchequer costs were about €266m, €10m and €1.3m for options A, B and C respectively.

approximately revenue neutral.¹³ Option (B), combining non-transferable bands with cuts in tax rates, gives rise to a sharp rise in married women's participation, and leaves men's participation almost unchanged. Option (C), using the revenue from restrictions on transferability to fund an increased child benefit, also boosts married women's participation, but leads to a fall in men's participation.

What about the total labour supply response, in terms of desired hours of work? Under option (A), the rise in average desired hours of work for women is almost offset by a fall in desired hours for men. Under option (B), which includes a significant cut in tax rates as well, the response of married women is more positive, and that of married men is less negative. As a result, the overall labour supply response for married couples is positive – and the response of single people, not simulated here in the present framework, would also be positive. Under option (C), the gain in tax revenue arising from non-transferability is applied to fund a rise in child benefit. This gives rise to a fall in male labour supply which is only partially offset by a rise in the labour supply of married women.

6. Conclusion

The research summarised here examined what labour supply response could be expected in Irish circumstances from married men and women to a number of forms of tax cuts and to structural reforms of the tax treatment of husband and wife. Results across the different forms of tax cuts (each with the same Exchequer cost) were quite similar in many respects, but the response of married women to a top rate tax cut or to band-widening was more than twice as strong as that of men, and more than twice as big as the women's response to a standard rate tax cut or allowance increase. Most of the change in desired hours appeared to be driven by changes in labour force participation.

Results on increased independence in the taxation of married couples showed that full individualisation of the standard rate tax band could have quite different impacts on labour supply, depending on the use made of the rise in tax revenue that would result. If this revenue were used to fund a general tax cut (through proportionate cuts in standard and top rates of tax) then men's participation would remain roughly constant, while married women's participation would rise by about 2½ percentage points. Average desired hours of work would rise by 0.4 hours per week. If, on the other hand, the revenue were used to finance an increase in child benefit, married women's participation would rise by more than the fall in married men's participation; but there would be a net fall in average desired hours of 0.2 hours per week.

Taken together, these results indicate that a revenue-neutral package involving individualisation of the tax bands and cuts in tax

¹³ As noted earlier, this is revenue neutrality on a static basis; increases (falls) in participation/hours would give rise to increased (reduced) revenues.

rates has more positive effects on labour supply than the more usual forms of tax cut – even if they cost IR£200m per annum more than the individualisation package. While the individualisation package (combining full individualisation with proportionate cuts in tax rates) has a greater labour supply impact than other tax cuts, the size of the effect needs to be put in context. Our results suggest a one-off rise of about 2½ percentage points in the labour market participation rate for married women in response to this policy package. This can be compared with an increase of about 30 percentage points in married women’s labour force participation since 1980.

More broadly, the results indicate that tax and welfare policy changes can have a significant impact on individual and family labour supply decisions. The size of the impact is an empirical question, best answered using methods of the type explored here. The new Irish evidence reviewed here indicates that simple “tax cutting” policies have less impact than restructuring of the tax treatment of couples, which acted to reduce the effective marginal tax rate on married women, one of the most responsive elements of the potential labour force. However, even this effect is modest compared to the trend rise in women’s participation driven by a range of factors, including rising real wage rates, higher education levels and social trends.

REFERENCES

- BLUNDELL, R., 2001. "Welfare reform for low income workers", *Oxford Economic Papers*, Vol. 53, No. 2, 189-214.
- BLUNDELL, R., and T. MACURDY, 1999. "Labor Supply: A Review of Alternative Approaches", in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics, Volume 3A*. Amsterdam: North-Holland.
- CALLAN, T., R. LAYTE, A. VAN SOEST, J. WALSH, 2003, *Taxes, Benefits and Labour Market Responses: New Evidence for Ireland*, Dublin: The Economic and Social Research Institute Policy Research Series Paper No. 48.
- CALLAN, T., M. KEENEY, B. NOLAN and J. WALSH, 2001. *Reforming Tax and Welfare*, Policy Research Series Paper No. 42 Dublin: The Economic and Social Research Institute.
- CALLAN, T. and A. VAN SOEST, 1996, "Family labour supply and taxes in Ireland", ESRI Working Paper No. 78.
- CALLAN, T. and B. FARRELL, 1991. Consultant's report in *Women's Participation in the Irish Labour Market*, Dublin: National Economic and Social Council, Report No. 91.
- DEPARTMENT OF FINANCE 1980. "Budget Speech 1980". Dublin: Stationery Office.
- DORIS, A., 2001. "The changing responsiveness of labour supply during the 1990s", in D. McCoy, D. Duffy, J. Hore and C. MacCoille (eds.), *Quarterly Economic Commentary* December.
- FAHEY, T., 1998, "Note on Labour Force and Family Status of Married Women, with reference to the implications of the tax treatment of married couples", Appendix 2 to Chapter 5, "Supporting families with their childcare responsibilities", in Commission on the Family, *Strengthening Families for Life*, Dublin: Stationery Office.
- KEANE, M. and R. MOFFITT, 1998, "A structural model of multiple welfare program participation and labor supply", *International Economic Review*, Vol. 39, pp. 553-589.
- KILLINGSWORTH, M. and J. HECKMAN, 1986, "Female labor supply: a survey", in O. Ashenfelter and R. Layard (eds.), *Handbook of Labor Economics*, Vol. 1, Amsterdam: North-Holland.
- O'DONOGHUE, C. and H. SUTHERLAND, 1999. "Accounting for the Family in European Income Tax Systems", *Cambridge Journal of Economics*, Vol. 23, No. 5, pp. 565-598.
- OECD, 1977. *The Treatment of Family Units in OECD Member Countries Under Tax and Transfer Systems*, A Report by the Committee on Fiscal Affairs, Paris: OECD.
- VAN SOEST, A., 1995, "Discrete choice models of family labor supply", *Journal of Human Resources*, Vol. 30, pp. 63-88.