MEMORANDUM

No 19/99

Progressive Taxes and the Labour Market – Is the Trade-Off between Equality and Efficiency Inevitable?

By
Knut Røed and Steinar Strøm

ISSN: 0801-1117

Department of Economics
University of Oslo

This series is published by the
University of Oslo
Department of Economics

P. O.Box 1095 Blindern
N-0317 OSLO Norway
Telephone: + 47 22855127
Fax: + 47 22855035
Internet: http://www.sv.uio.no/sosoek/
e-mail: econdep@econ.uio.no

In co-operation with
The Frisch Centre for Economic Research

Gaustadalleén 21
N-0371 OSLO Norway
Telephone: +47 22 95 88 20
Fax: +47 22 95 88 25
Internet: http://www.frisch.uio.no/
e-mail: frisch@frisch.uio.no

List of the last 10 Memoranda:

No 09	By Sverre A.C. Kittelsen: Monte Carlo Simulations of DEA Efficiency Measures and Hypothesis Tests. 61 p.
No 10	By Asbjørn Rødseth and Ragnar Nymoen: Nordic wage formation and unemployment seven years later. 52 p.
No 11	By Finn R. Førsund and Kjell Ove Kalhagen: Efficiency and Productivity of Norwegian Colleges. 38 p.
No 12	By Steinar Holden: Wage setting under different Monetary regimes. 40 p.
No 13	By Tor Jakob Klette, Jarle Møen and Zvi Griliches: Do subsidies to commercial R&D reducemarket failures? Microeconometric evaluation studies. 36 p.
No 14	By Knut Røed and Tao Zhang:Unemployment Duration in a Non-Stationary Macroeconomic Environment. 32 p.
No 15	By T. Hægeland, T.J. Klette and K.G. Salvanes: Declining returns to education in Norway? Comparing estimates across chorts, sectors and overtime. 21 p.
No 16	By Oddbjørn Raaum and Tom Erik Aabø: The effect of schooling on earnings: The role of family background studied by a large sample of Norwegian twins. 31 p.
No 17	By Erik Hernæs, Marte Sollie and Steinar Strøm: Early retirement and Economic Incentives. 27 p.
No 18	By Erik Hernæs: Fewer in Number but Harder to Employ: Incidence and Duration of Unemployment in an Economic Upswing. 28 p.

A complete list of this memo-series is available in a PDF® format at:
http://www..sv.uio.no/sosoek/memo/

22 July 1999

Progressive Taxes and the Labour Market – Is the Trade-Off between Equality and Efficiency Inevitable?

By Knut Røed[†] and Steinar Strøm[‡*]

Abstract

Does an income tax harm economic efficiency more the more progressive it is? Public economics provides a strong case for a definite 'yes'. But at least three forces may pull in the other direction. First, low-wage workers may on average have more elastic labour supply schedules than high-wage workers, in which case progressive taxes contribute to a more efficient allocation of the total tax burden. Second, in non-competitive labour markets, progressive taxes typically encourage wage moderation, and hence reduce the equilibrium level of unemployment. And third, if wage setters have egalitarian objectives, progressive taxes may reduce the need for redistribution in pre-tax wages, and hence increase the demand for low-skilled workers. This paper surveys the theoretical, as well as the empirical literature about labour supply, taxes and wage setting. We conclude that in a second best world, the trade-off between equality and efficiency is not always inevitable.

JEL Classification: H24, J30, J58

Keywords: Tax progressivity, dead-weight loss, redistribution

---

[†] The Ragnar Frisch Centre for Economic Research, Oslo.

[‡] The Department of Economics, University of Oslo and The Ragnar Frisch Centre for Economic Research, Oslo.

[*] This paper is part of the project "Income taxes, wage formation and distribution" at The Ragnar Frisch Centre for Economic Research, Oslo. Financial support from the Research Council of Norway is gratefully acknowledged. We wish to thank Vidar Christiansen, Jon Strand and Fredrik Wulfsberg for valuable comments. Correspondence to Knut Røed, The Ragnar Frisch Centre for Economic Research, Gaustadalleen 21, 0349, Oslo, Norway. E-mail: knut.roed@frisch.uio.no.

1 Introduction

The tax reforms that took place in many industrialised countries during the 1980's and early 1990's involved a sharp movement away from highly progressive income tax schedules and towards broader tax bases. In some countries, the marginal tax rates faced by top rate earners dropped from 70-80 per cent to around 40-50 per cent (Blundell, 1996). These reforms emanated from a growing concern about the costs of redistribution. Their aim was to reduce the perceived strong disincentives to work, and hence enhance economic efficiency. Their intellectual basis was provided by public economics.

Economic theory typically pictures the trade-off between equality and efficiency as inescapable. In one of the most popular textbooks about public economics (Stiglitz, 1986, p. 481), we learn that it is possible to redistribute income from the rich to the poor through progressive taxation, but are warned that we do so at the expense of a loss in economic efficiency: "The government can reduce the degree of inequality but only at the expense of a larger dead-weight loss." Empirical research, particularly related to the evaluation of labour supply decisions, has also indicated that the efficiency loss associated with high marginal tax rates is far from negligible. And in standard models of labour supply, it is the marginal, and not the average, tax rate that induces waste of resources. Hence a progressive tax system, which by construction may involve high marginal, relative to average, tax rates, also involves a large efficiency loss relative to the revenues collected.

More recently however, doubts have been raised about the unambiguous efficiency loss associated with tax progressivity. Two mechanisms have been focused upon. First, an almost inescapable by-product of revenue neutral tax reforms that reduce tax progressivity, is that for some very low wages, both average and marginal tax rates increase (OECD, 1995, p. 42; Leibfritz et. al., 1997, p. 64). There is some evidence suggesting these low-wage workers are subject to a particularly elastic labour demand (OECD, 1995, p. 86), as well as labour supply (Juhn et al, 1991, p. 113). Hence, there is a possibility that a reduction in tax progressivity shifts the tax burden away from the less responsive workers and towards the more responsive workers, thereby in-

creasing rather than reducing efficiency costs. Secondly, a reduction in tax progressivity may affect the wage setting behaviour, and hence also change the (not necessarily efficient) equilibrium rate of unemployment. As it turns out, models of imperfectly competitive wage setting typically predict a negative relationship between tax progressivity and equilibrium unemployment (see e.g. Sørensen, 1997, p. 227).

There are a number of problems associated with the way tax progressivity is analysed in the labour supply-, as well as in the wage setting literature. Even though the optimal tax literature following Mirrlees (1971) does provide important theoretical insight about the relationship between worker heterogeneity and the property of efficient income tax schedules (see e.g. Atkinson and Stiglitz, 1980), little empirical work has been done in order to the characterise this heterogeneity and to assess its policy implications. The representative agent model prevails as the dominating foundation for applied work. This approach has served to isolate the substitution effect in labour supply, as well as in wage setting models. But at the same time it has made it difficult to analyse the implications of changes in the tax system that really redistributes the burden of taxation. The fact that the income tax is only one element of a larger welfare support system that, particularly for low paid workers, strongly affects the participation decision is often overlooked. And while one branch of the literature focuses on labour supply decisions, given the wage rate, and another branch focuses on wage determination, given the labour supply, there have been few attempts to integrate these approaches into a common framework.

The purpose of this survey is to bring the various strands of the literature together, to add some more realistic assumptions about worker heterogeneity, and then to evaluate what economics really has to say about the trade-off between equality and efficiency in the construction of income tax policies. We limit ourselves to a static view of taxation; hence we do not discuss how the tax system may affect investment in human capital or allocation of effort and income over time. We also sidestep the important issues of tax avoidance and tax evasion, although such activities obviously may be related to the degree of progressivity. The next section discusses the concept of tax progressivity. We point out that the notion of progressivity embedded in the labour supply- and wage setting literature is very different from that embedded in the income distribution literature. Section 3 discusses the relationship between tax progressivity

and labour supply. We first illustrate why the degree of progressivity, from a theoretical point of view, has ambiguous effects on economic efficiency, and then evaluate the empirical evidence. Section 4 discusses the relationship between tax progressivity, wage determination and equilibrium unemployment. Again, we start out with the theoretical arguments before we take a look at empirical evidence. A key point in our evaluation of the literature is that there are some important missing links between the theoretical arguments and the associated empirical evidence, primarily related to the neglect of worker-, as well as union, heterogeneity. Section 5 offers some concluding remarks. Unsurprisingly perhaps, we conclude that there is no such thing as a unique trade-off between equality and efficiency. Instead, there are different trade-offs associated with different tax policies. For some redistributive tax policies, there is probably no trade-off at all; i.e. the two aims of equality and efficiency are not in conflict.

2 The Concept of Tax Progressivity

How do we measure the degree of progressivity embedded in a tax system? A progressive tax is usually defined as a tax which is levied at a rate that increases as the tax base increases. Thus, a tax is progressive when the marginal tax rate exceeds the average tax rate, while it is regressive when the average tax rate exceeds the marginal tax rate. If the average tax rate increases with income for all income levels, the tax *system* is said to be *uniformly progressive*. Of course, in progressive tax systems there has to be an upper bound for the level of the average tax rate. This upper bound is defined by the highest marginal tax rate, and the average tax rate approaches this highest marginal tax rate when income goes to infinity. In most countries the tax systems are *not* uniformly progressive, and the degree of progressivity varies substantially with the level of personal- or family income.

2.1 Local Measures of Tax Progressivity

Accordingly, one may construct *local* measures of progressivity; i.e. progressivity schedules that for each possible income level measure the degree to which the average tax rate increases when the income increases. Two of the still most popular measures were proposed by Musgrave and Thin (1948). The first, *tax liability progression* (TLP), is the elasticity of the tax bill with respect to the pre-tax income. The second,

residual income progression (RIP), is the elasticity of the post-tax income with respect to the pre-tax income. Let *y* denote the pre-tax income, and let *t(y)* be the tax paid as a function of that income. The two local tax progressivity measures, *TLP* and *RIP*, are related to the marginal and the average tax rate in the following simple way:

$$TLP(y) = \frac{t'(y)}{\bar{t}(y)}, \quad \bar{t}(y) > 0,$$

$$RIP(y) = \frac{1 - t'(y)}{1 - \bar{t}(y)} \quad \bar{t}(y) < 1,$$

where *t'(y)* is the marginal tax rate and $\bar{t}(y)$ is the average tax rate associated with the income level *y*. A flat positive tax yields *TLP=RIP=1*, no matter the tax rate and no matter the income. A progressive tax yields *TLP>1* and *RIP<1* at least for some levels of income, while a regressive tax yields *TLP<1* and *RIP>1*. In this paper, we focus primarily on the *RIP*-measure. This measure has the advantage of being defined even in income intervals with a zero tax rate. Moreover, it turns out to play an important role in models of wage setting. Note that the more progressive the tax is, the *lower* is RIP.

2.2 Global Measures of Tax Progressivity

As workers have different incomes, and hence are subject to different degrees of progressivity, we still do not have the tools needed to compare the degree of progressivity embedded in different tax *systems*. One tax system may yield higher progressivity than another system for some workers, but lower progressivity for other workers. This problem does not arise as long as one operates within the framework of a *representative agent model*. In such a model, all workers are assumed equal; i.e. they have the same income and they are subject to the same degree of tax progressivity. But if all workers were equal, the whole issue of tax progressivity would be irrelevant. There would be no reason to impose distortionary income taxes, at least not in a competitive economy (a lump-sum tax could raise the same money at a zero efficiency cost). On the other hand, if we take into account the fact that workers have different wage opportunities and different preferences with respect to consumption and leisure, it is very difficult to rank different tax systems according to some simple measure of progressivity. There is no generally accepted scalar measure of *global* tax progressivity. A progressivity ranking of tax systems involves subjective judgements that are very

similar to the judgements needed to construct scalar measures of inequality in general. While local measures of tax progressivity focus on the tax schedule faced by a singel (representative) agent, global measures focus on the degree of redistribution that actually takes place between agents, i.e. the redistributive effect. For example, the Reynolds-Smolensky (1977) measure of global progressivity is the difference between the pre- and post tax Gini coefficients, for a given income distribution (see Creedy, 1996, Chapter 3, for an overview of various progressivity measures). This measure does of course not only depend on the tax system, but also on the actual distribution of incomes. It should be noted that to the extent that taxes affect behaviour, the distribution of pre-tax income is affected by the design of the tax system. To assess the impact on the post-tax distribution of changes in taxes thus requires that the impact on the pre-tax distribution is accounted for.

Gini based measures of global tax progressivity are not directly related to the local measures. Ebert (1992) proposed a global measure that does build on local tax progressivity; i.e. the geometric mean of the RIP's faced by all the workers in the economy. But local measures, such as the RIP, do not really capture the degree in which the tax burden is redistributed *between individuals* with different incomes; they only reflect the relationship between the average and the marginal tax rates at each given point in the income distribution[1]. One way to see this is to consider a (purely hypothetical) system of two different flat tax rates; one low rate assigned to people with a low income potential (determined by nature) and one high rate assigned to people with a high income potential. As both these groups face a flat tax rate, there is no local progressivity, no matter the income level at which is it measured. Nevertheless, the system entails a strong element of redistribution (that would be reflected in the Gini-based global measures). As becomes clear in the next section, this element of pure redistribution may have important efficiency implications, even though it is not reflected in local measures of progressivity.

[1] Jakobsson (1976) showed, however, that RIP in contrast to e.g. TLP, satisfies the natural requirement that if one tax system is everywhere, according to the local measure, more progressive than another, then it should also be unambiguously more redistributive than the other (in terms of the Lorenz criterion).

3 Tax Progressivity, Labour Supply and the Dead-weight Loss

A tax is *non-distortionary* if agents cannot affect the amount of taxes paid through their own behaviour. Thus, a non-distortionary tax is independent of individual income and wealth. Such a tax is often referred to as a *lump-sum tax*.

Non-distortionary, as well as distortionary, taxes do of course affect behaviour. Indeed, the whole point of imposing a tax is to affect behaviour, in the sense that agents must either work more or consume less in order to make resources available for public purposes. In general, a tax affects behaviour through two channels: First, it makes the taxpayer poorer, and this typically causes a fall in the consumption of goods as well as leisure (the *income effect*). Second, the tax may affect the reward (cost) attached to working more (less) hours (the *substitution effect*). The former of these effects applies to all taxes, including lump-sum taxes. The latter applies to distortionary taxes only. A tax on labour-income typically causes a worker to work more hours because he is worse off (and therefore reduces the consumption of all normal goods, including leisure) and fewer hours because the take home wage rate is reduced. It is far from obvious which of these two effects that dominates the labour supply decision. But from a pure efficiency point of view, it doesn't really matter. In the dead-weight literature, the size of the efficiency loss is determined by the substitution effect alone.

3.1 The Dead-Weight Loss

The efficiency loss arising from distortionary taxation has many labels in the literature, the most common being the *dead-weight loss* and the *excess burden*. Since the loss arises due to the substitution effect only, the most common measures of the dead-weight loss (excess burden) are based upon hypothetical side payments that 'remove' the income effect. Two alternative dead-weight measures, due to Hicks (1939), still dominate the literature: The first of these measures is based on the concept of *compensating variation* (CV). Traditionally its starting point has been an agent's indirect utility in the absence of an existing tax, or under a different tax regime (the proposed reform). The compensating variation is the side payment required to keep the utility level fixed when the proposed reform system (or the situation without tax) is replaced by the currently existing tax system, i.e. it is a money metric measure of the welfare effects associated with a tax reform (or the abolishment of a tax), using the hypotheti-

cal reform-alternative (or the situation without tax) as the reference case. The dead-weight loss attached to a particular tax imposed on a particular person is given by the difference between the size of the required side payment and the (additional) revenue raised by the current tax. The other measure builds on *equivalent variation* (EV). The equivalent variation is the amount that an agent is willing to pay in order to get the proposed reform effectuated (or an existing tax abolished) i.e. it is a money metric measure of the welfare effects associated with a tax reform, using the prevailing situation as the reference case. The dead-weight loss is again measured as the difference between the required side payment and the revenue raised by the tax.

While the CV- and EV measures both have clear-cut theoretical interpretations; the EV-measure is typically preferred in empirical settings. The reason is that it has the advantage of using existing prices and wages (and not those that would apply under a different tax regime) as the basis for comparison, thereby facilitating comparison of several alternative policy changes, see Kay (1980). As long as one focuses on small changes in the tax system, the two measures typically yield the same result. Welfare assessments of more radical tax reforms, such as the abolishment of all distortionary income taxes, are much more controversial, as the models used for such assessments are estimated on data generated in environments that resemble the current tax system much more than the alternative. One can question whether estimated models are valid at all under purely hypothetical circumstances that are completely different from the circumstances under which the estimates are obtained. Hence, a better approach may be to use a flat tax as a reference case in order to evaluate the dead-weight loss associated with existing progressive taxes. The advantage with this procedure is that it does not deviate too much from the data generating process, while at the same time the marginal price of leisure, given a common wage, is equated across individuals, just like in the no-tax-world (in the absence of other imperfections).

3.2 The Representative Agent

We first show why a progressive tax yields such a large dead-weight loss, compared to a flat tax. To keep the argument as simple as possible, we use the tax-free world as a reference, and hence build the discussion on the CV-measure. Consider an agent who is free to determine the number of hours worked, h, at any given wage rate w.

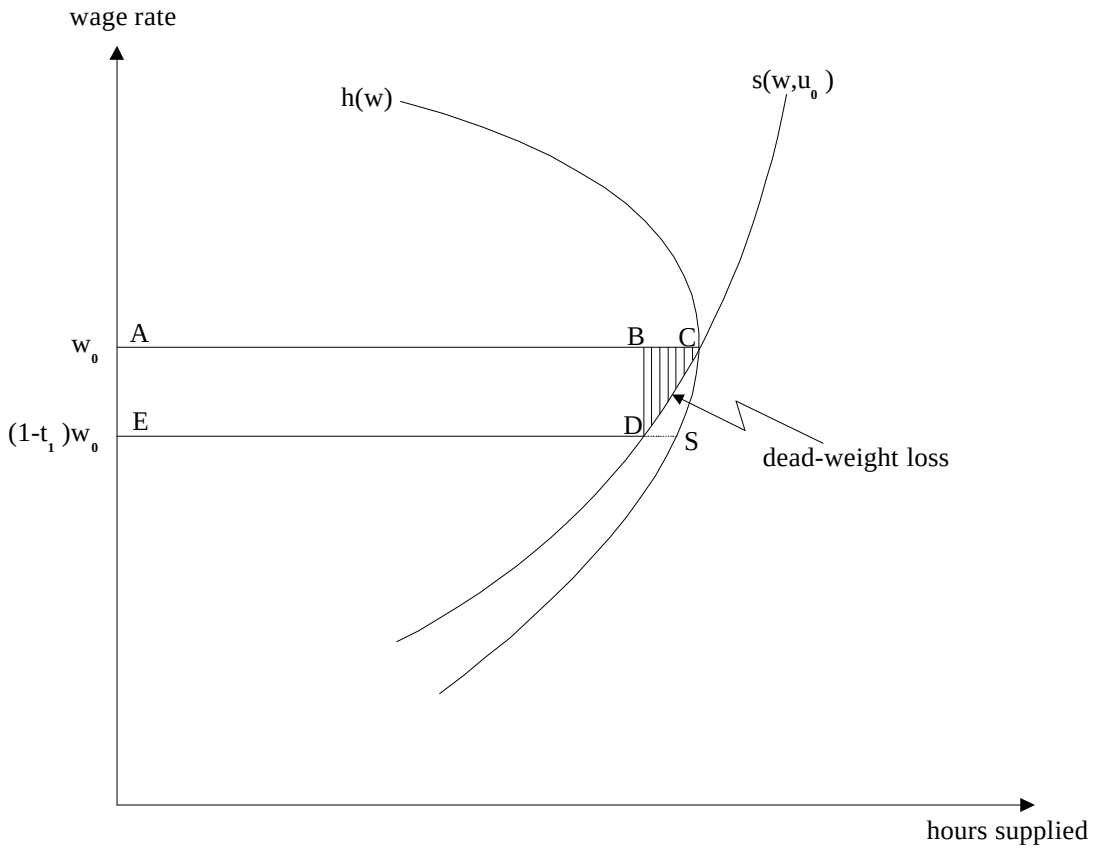

Figure 1. Dead-weight loss for a representative agent under a flat tax regime

Figure 1 depicts how the number of hours worked may depend on the wage rate through the backward bending labour supply curve *h(w)*. Assume that this worker faces the wage w_0 and a flat tax rate t_1. Now, in the absence of that tax, the worker would have settled in point *C* and enjoyed the utility level u_0. As we focus on compensated variations from this hypothetical starting point, we draw the compensated labour supply curve $s(w,u_0)$ through point *C*. This curve depicts the labour supply that would have occurred if any change in the take home wage rate (from w_0) were accompanied by a side payment that maintained the utility level fixed at u_0. In the flat tax regime, the take home wage rate for our representative worker is given by $(1-t_1)w_0$. How large is the compensation needed to keep the utility unchanged at the level associated with no tax at all? Consider first the case in which the tax is infinitesimally small and has only a negligible effect on behaviour. The required compensation would then approximately equal the size of the tax per hour, multiplied with the number of hours worked (the line segment AC in the figure). But as the tax becomes larger, the required additional compensation declines, as the hours worked also de-

clines. Hence, it should not be too difficult to realise that the total compensation required (CV) is equal to the area ABCDE. Now, if such a compensation were really paid, our agent would have settled in point D, hence the total tax bill would have been equal to ABDE, which is less than the compensation required. Hence, even if all tax revenues were paid back to the worker in the form of a lump-sum compensation, this would not be enough to maintain the worker's utility level. The difference between the compensation required and the tax revenues collected is the dead-weight loss (the

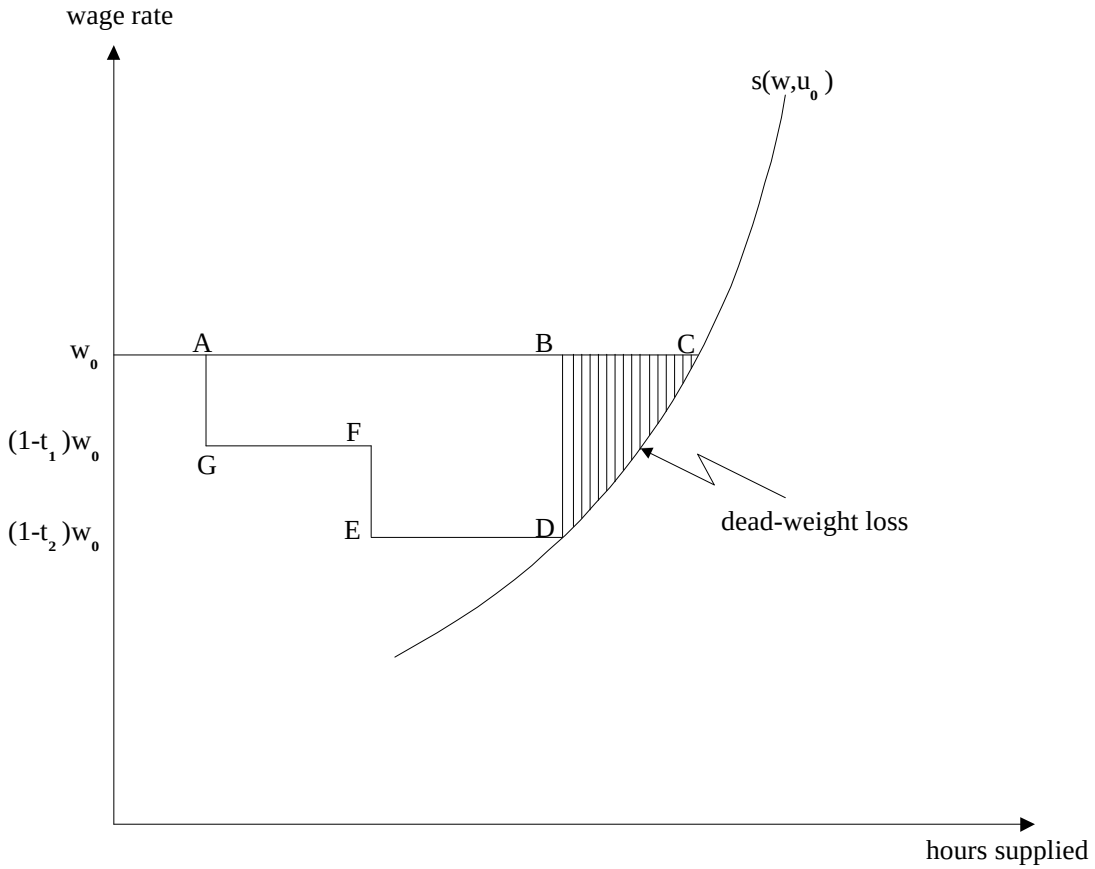

Figure 2. Dead-weight loss for a representative agent under a progressive tax regime

hatched area in the figure). These are the resources that may be considered completely lost through the distortionary taxation. Measurement of the dead-weight loss according to the EV measure is very similar. The difference is that we would have drawn the compensated labour supply curve through the point applying to the situation with the tax in place (point S), and used that curve as the tool for the analysis.

Figure 2 illustrates how the dead-weight loss increases sharply when we replace the flat tax (t_1) with a progressive tax with three different rates ($0<t_1<t_2$), that raises the

same compensated revenue from our representative agent. The important point to note is that, with respect to the dead-weight loss, it is only the tax rate that applies at the point of intersection of the after tax wage schedule and the compensated labour supply curve – *the marginal tax rate* - that matters. Hence, the dead-weight loss, measured in this way, would have been exactly the same had the highest tax rate (t_2) been applied to all income. The tax rebate given to the first part of the income works just like a lump-sum transfer, or, in the words of Burtless and Hausman (1978), a 'virtual income'. It does not buy any reduction at all in the dead-weight loss[2].

This argument constitutes an important element of the scientific foundation for structural reforms that aim at cutting the highest marginal rates of taxation. It implies that a *given tax payment* collected from a *given individual* unambiguously imposes a larger burden on that individual, the higher is the marginal rate of taxation. But the argument is built on a platform of strong assumptions. First, it is assumed that workers are free to work the number of hours they themselves prefer. Second, it is assumed that wage rates are not affected by the tax system. And third, to the extent that the argument is given an economy-wide interpretation, it is assumed that all workers are equal. Empirically, these assumptions are of course utterly irrelevant. Nevertheless, most economist would probably argue that they serve the useful purpose of providing a simple framework for discussing efficiency costs related to taxes. The typical view is that relaxation of one or more assumptions may complicate things and dilute some very strong conclusions, but not alter the basic arguments. This may be wrong. There is a possibility that more realistic assumptions will undercut the premises for the presumed inescapable trade-off between equality and efficiency. That is the issue to which we now turn.

3.3 Heterogeneous Workers and the Aggregation Problem

For expository reasons we proceed by assuming that agents are free to choose any hours of work they prefer and that wages are determined independently of the tax system (we return to these issues in later sections), but we revoke the assumption that

[2] Note that the virtual income effect causes a shift in the uncompensated labour supply curve (not shown).

all workers are equal; i.e. we abandon the representative agent framework. Workers may have different productivities, and hence face different wage opportunity sets, and/or they may have different preferences. The first thing to note is that in a heterogeneous world, it is not possible to separate equality- from efficiency considerations in any impartial way. The reason is that, even though we (in principle) can measure the dead-weight loss that accrues to each person, we cannot aggregate these losses without further assumptions that implicitly place weights on each individual. Moreover to compare the welfare losses of some individuals to the welfare gains of others we need cardinal utilities that are comparable across all individuals. Provided that we are ready to place political weights on each individual and employ comparable cardinal utilities, aggregate welfare gains can be assessed. If the social welfare gain caused by a given increase in a person's utility is independent on whether the increase is allocated to a person with high or low utility, the social welfare function is said to be *utilitarian*. In that case, the aggregate money-metric of the welfare improvement associated with e.g. a tax reform is given by the sum of individual utility gains – as measured by compensating or equivalent variations. However, if one is not ready to embrace the utilitarian welfare function, egalitarian considerations may imply that a tax reform is not deemed desirable from a social point of view, even when there is a net positive willingness to pay for the reform in the population. The reason is that the winners are not in general able to compensate the losers, as lump-sum transfers are not feasible (otherwise it wouldn't be necessary to impose distortionary taxes in the first place). If we limit the analysis to ordinal utilities, welfare gains and losses for different persons cannot be compared in any meaningful way, and welfare analyses are limited to a simple identification of winners and losers. In that case an aggregate measure that can be reported is the fraction of losers and winners - measured by the sign of the compensating or equivalent variation - from a tax reform.

The sum of compensating or equivalent variations may alternatively be given a less normative interpretation. For example, the extra dead-weight loss (the efficiency cost) associated with progressive (rather than flat) taxation may be measured as the total willingness to pay for a flat-tax reform (keeping total tax revenues constant) in the population. The presumed social gain, in terms of improved distribution may be associated with particular properties of the joint distribution of individual gains and losses, and household characteristics (Hammond, 1990). This implies that the sum of com-

pensating or equivalent variations is associated with *efficiency*, rather than *welfare*, while their distribution is used to asses welfare changes related to the degree of *equality*. There is a net efficiency cost associated with a given tax reform if the sum of compensating (or equivalent) variations is negative. This forms the basis for a popular efficiency cost measure, the *marginal efficiency cost of redistribution* (MECR) (Ballard, 1988). Consider a very small, proposed change in the tax system that redistributes income from the rich to the poor. For each individual, the CV is the lump-sum transfer required under this alternative tax system in order to be indifferent to a restoration of the existing system. Hence, if a CV is positive, the proposed change improves welfare. Let $\sum CV_+$ be the sum of CV's for the (poor) winners and let $\sum CV_-$ be the sum of CV's for the (rich) losers. The marginal efficiency cost of redistribution associated with that particular proposal is then given by:

$$MECR = \frac{-\sum CV_-}{\sum CV_+} - 1$$

Hence, if the rich lose more than the poor gain, there is a positive efficiency cost associated with the proposal, and vice versa.

Consider again Figure 1. It is obvious that the size of the dead-weight loss for each individual depends on two factors: The marginal tax rate and the slope of the compensated labour supply curve. A given marginal tax rate entails a larger dead-weight loss the larger is the compensated elasticity of labour supply. If workers have different compensated elasticities of labour supply, the aggregate dead-weight loss cannot be minimised by a common flat-tax rate. Let us, for the sake of argument, assume that workers with low wages on average have a more elastic (compensated) labour supply schedules than workers with high wages have (the empirical relevance of this assumption is discussed in the next section). If this is true, it provides, ceteris paribus, a strong efficiency argument for taxing the poor low-wage workers less than the rich high-wage workers, i.e. a progressive tax system. To illuminate this point, let us replace our representative agent with the simplest possible alternative, namely two representative agents; one with a high and one with a low wage rate. For simplicity, we assume that the two agents have identical preferences, hence they are represented by a common labour supply curve. But since they have access to different (fixed) wage rates (due to different productivities) they adapt differently. Figure 3 illustrates how a

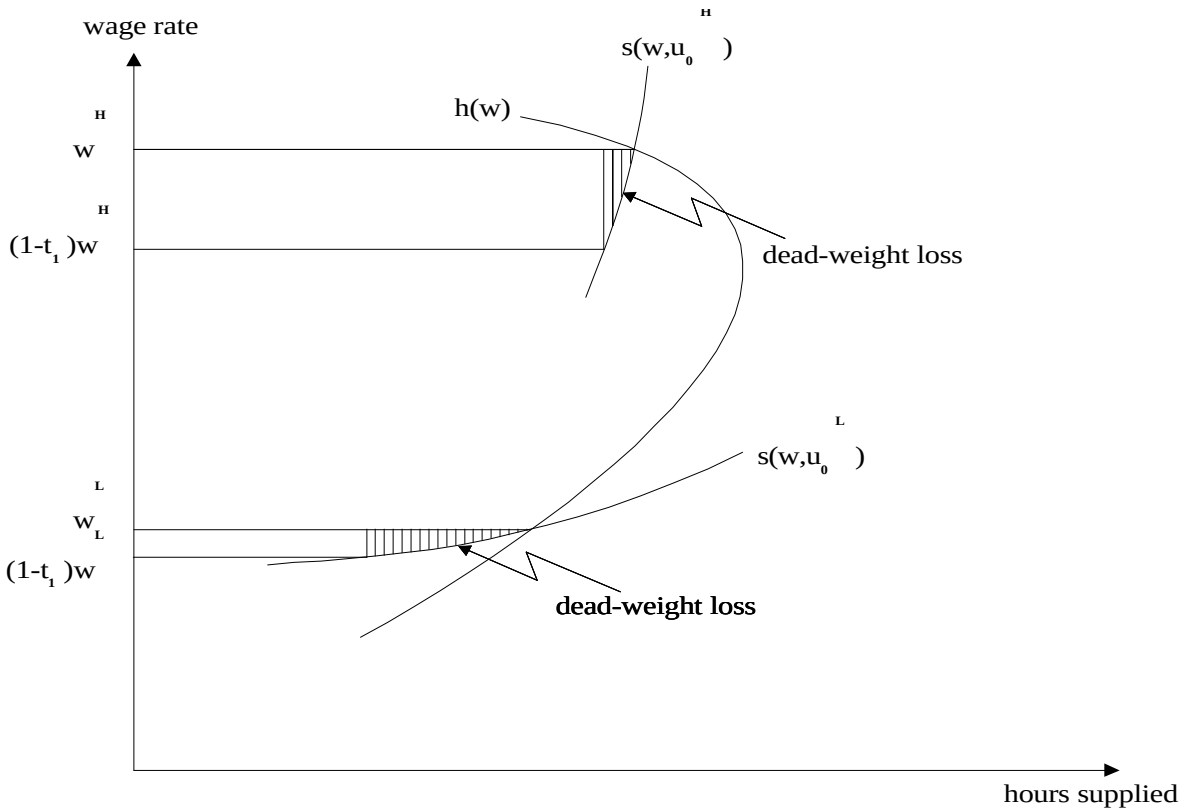

Figure 3. Dead-weight losses for two representative agents under a flat tax regime

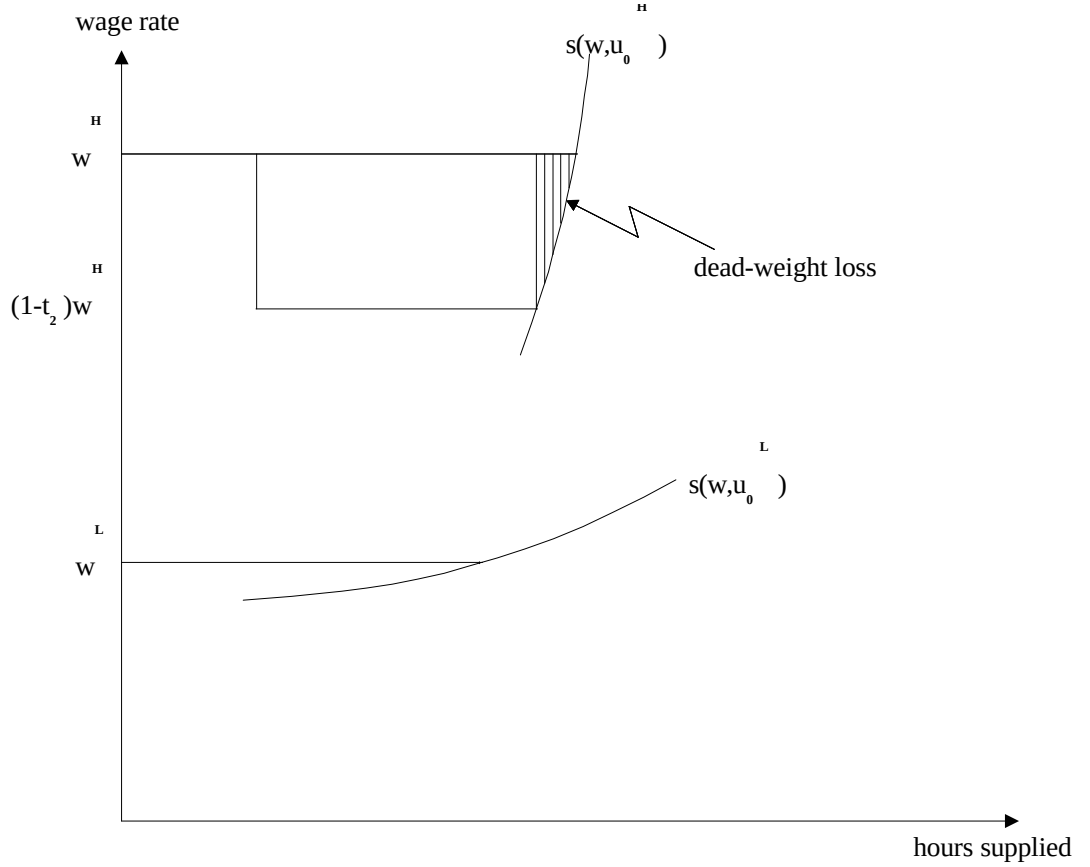

Figure 4. Dead-weight losses for two representative agents under a progressive tax regime

flat-tax system may impose a larger dead-weight loss on the person with the low wage (w^L) than on the person with the high wage (w^H). And figure 4 illustrates that a progressive tax system (in this case represented by a basic income tax allowance, such that the sum of compensated tax revenues is unchanged) may well reduce the aggregate dead-weight loss. This happens even though the local measure of progressivity either remains unchanged (for the poor) or displays increased progressivity (for the rich). The key to understanding this result is that the progressive tax system transfers the burden of taxation from a person at an income level associated with a very elastic labour supply, to a person at an income level associated with a very inelastic labour supply. Hence, with respect to the aggregate dead-weight loss there are two conflicting effects at work. For each person it is true that the tax system yields a larger dead-weight loss the higher is the degree of local progressivity, *given the tax bill paid by that person*. On the other hand, the higher local progressivity is a by-product of a tax system that allocates the tax bills *between* individuals in a more efficient way.

Now, it is well known that the compensated labour supply analysis may be inappropriate as a tool for comparing different tax systems that are all distortionary. The reason is that in the presence of e.g. a coexisting income tax, even a lump-sum tax involves efficiency considerations, through its 'revenue effect'. Given that leisure is a normal good, a distortionary income tax implies (ceteris paribus) that 'too much' leisure is consumed. An additional lump-sum tax mitigates this distortion, because its pure income effect discourages the 'excessive' consumption of leisure. Accordingly, each extra Euro raised through the lump-sum tax raises more than one Euro in total taxes. This implies that the income elasticity of labour supply is also relevant for efficiency considerations, and many economists prefer to use uncompensated – rather than (or in addition to) compensated – elasticities to evaluate efficiency costs associated with particular taxes.

To sum up, we do not claim that income redistribution through the tax system is a socially costless endeavour. The point is that there may be conflicting effects at work, even in a competitive labour supply model, such that more progressivity does *not always* imply more waste of resources. Even if the social welfare function places no weight at all on equality, it may be theoretically possible to calculate an optimal income tax schedule (in the second best sense) that involves an element of progressiv-

ity. Nevertheless, it may very well be the case that today's existing tax systems are too progressive, or that they display the wrong type of progressivity.

3.4 Empirical Evidence

Now, are these examples of efficiency-improving tax progressivity empirically relevant? Available evidence indicates that there is an enormous variation in (compensated as well as uncompensated) elasticities of labour supply. The literature typically focuses on three demographic groups with different labour supply behaviour: men, married women, and lone parents (mothers). The labour supply of male workers is usually found to be very inelastic on average, while the labour supply of married women (working short hours or not participating at all in the labour market) and lone parents is found to be very elastic (see. e.g. Blundell, 1997, for a recent overview).

There is little direct evidence regarding the relationship between labour supply elasticities and wage rates. Indeed, the sign of the derivative of the labour supply elasticity with respect to the wage rate is more often imposed than investigated. Popular labour supply specifications, such as the linear and the log-linear, restrict this derivative (ceteris paribus) to be positive or zero, respectively[3]. As pointed out by Triest (1994, p. 149), evaluations of tax reforms that build on linear labour supply models (restricting the wage elasticity to rise with the wage level, and decline with hours worked) are seriously biased *against* reforms that increase marginal tax rates for the rich (when the rich are rich because they have high wage rates). Table 1 offers an overview of the existing evidence that we believe come closest to throwing some light on the issue. The limited evidence that does exist indicates that when a flexible functional form of the utility function is applied - such that the empirical evidence is allowed to determine how the wage elasticity varies with income – one finds that these elasticities are declining with household income

[3] For example, the well-known Hausman approach to labour supply estimation applies a linear supply curve (Hausman, 1979; 1980; 1981; 1985a; 1985b; Hausman and Ruud, 1984; Blomquist, 1983).

Table 1
Empirical evidence regarding the relationship between labour supply elasticities and (household) income or wages.

Authors	Coverage	Methodological approach/type of data	Results
Devanzo et al (1973)	United States. Men.	Labour supply model, estimated on micro data, including information about participation/non-participation and hours.	Virtually all of the labour supply wage- and income responsiveness is found at or near the zero-hours point.
Borjas and Heckman (1978)	United States. Men.	Labour supply model, estimated on micro data, including information about participation/non-participation and hours.	Labour supply estimates are more responsive to wages and incomes when participation decisions are accounted for than when only hours of work, given participation, are used in estimating labour supply.
Arrufat and Zabalza (1986)	United Kingdom, 1974. Married women.	Micro data based on the General Household Survey. Labour supply model, with husbands' labour supply treated as exogenous.	The estimated total labour supply elasticity for married women is 2.03, out of which 1.41 is driven by participation decisions.
Dagsvik et al (1988)	France, Women, 1979	Labour supply model, estimated on micro data collected from the INSEE survey 'Budgets des Familles 1978-1979'. The data include information about participation/non-participation and hours.	The estimated total labour supply elasticity is on average around 3, out of which approximately 1.4 is driven by participation decisions.
Juhn et al (1991)	United States, 1970-89. Men.	Current population survey (CPS) data. Fraction of year spent working regressed on individual wage rates (or estimated wage rates).	The participation decision is more elastic for workers with low wages (or low potential wages). For example, the estimated partial labour supply elasticities are approximately five times higher for workers in the 1-10 percentile than for workers in the 61-100 percentile of the wage distribution.
Aaberge et al (1995)	Norway, 1979. Married couples.	Labour supply matching model, estimated on micro data collected from the Level of Living Sample Survey.	Both participation and hours elasticities are higher the lower is household income. For example, for all men (women) the estimated uncompensated labour supply elasticity is 0.45 (1.82) out of which 0.29(0.83) is due to participation. For the 10 per cent poorest, the corresponding numbers are 2.23(3.09) and 1.89 (1.85).
Aaberge et al (1999b)	Married couples in Italy (1987) Norway (1986) and Sweden (1981)	Participation decisions as well as hours of work (for Sweden only working couples). Accounts for non-convex budget sets and restrictions on hours offered in the market.	For Italy see next entry. For Norway the simulated direct uncompensated labour supply elasticities for all men (women) are 0.28 (0.91) out of which 0.17 (0.37) is due to participation. For working couples in Sweden in 1981 the simulated direct labour supply elasticities are -0.02 for men and 0.07 for women.

Table 1
Empirical evidence regarding the relationship between labour supply elasticities and (household) income or wages.

Authors	Coverage	Methodological approach/type of data	Results
Aaberge et al (1999a)	Italy, 1987. Married couples.	Labour supply matching model estimated on data from the Survey of Household Income and Wealth.	The simulated uncompensated direct elasticities for men (women) are 0.05 (0.74) out of which 0.04(0.65) is due to participation. For the 10 per cent poorest the corresponding numbers are 0.08 (3.44) and 0.05 (2.84).
Moffit and Wilhelm (1998)	United States, 1983-1989. Affluent men.	Data is collected from the Survey of Consumer Finances (SFC) and used to evaluate hours of work responses to the 1986 Tax Reform Act.	The labour supply of high-income men is inelastic with respect to the marginal tax rate. There is no evidence of changes in hours of work in response to the marginal tax rate reductions legislated in the 1986 Tax Reform Act.

A robust finding in the literature is that the participation elasticity is much higher than the elasticity of hours worked, given participation. One of the first to point out that virtually all of the wage- and income responsiveness in male labour supply is concentrated at or near the zero-hours point was Devanzo et al (1973). In particular among married women, a recurring result in the empirical labour supply literature is that economic incentives have a much stronger impact on the decision to participate in the labour force than on hours worked (Cogan, 1981; Heckman, 1978; Mroz, 1987).

In the more recent approaches to labour supply (Hausman and Ruud, 1984; Ilmakunnas and Pudney, 1990; Dickens and Lundberg, 1993; Aaberge et al, 1995; 1999a; 1999b) the participation issue associated with married couples is taken account of through the formulation of a joint decisions problem, given a non-convex budget constraint. The family labour supply studies indicate that the labour supply of married women is extremely sensitive to economic incentives and that this sensitivity to a large extent is accounted for by the discrete participation decisions[4].

One of the few papers that report elasticities for different income deciles, and also decompose these elasticities in terms of participation and hours of work decisions (given participation) is Aaberge et al (1999b). These elasticities are derived from simulations

[4] Similar results are obtained for lone mothers, see e.g. Walker (1990) and Rønsen and Strøm (1991).

on labour supply models estimated on micro-data from Italy and Norway. The utility function is a Box-Cox transformation of leisure for wife and husband and of household consumption. In estimating the models most of the details in the tax and transfer structure are accounted for. The reported elasticities, which we reproduce in Table 2, are the average over the persons in the relevant deciles of the household income distribution. The table includes own wage elasticities as well as elasticities with respect to the spouses' wage (cross wage).

Table 2.
Uncompensated aggregate labour supply elasticities, Italy 1987 and Norway 1986

	Male elasticities				Female elasticities			
	own wage		cross wage		own wage		cross wage	
	Italy	Norway	Italy	Norway	Italy	Norway	Italy	Norway
Elasticity of the probability of participation, all households	0.046 (0.001)	0.17 (0.004)	-0.081 (0.002)	-0.03 (0.004)	0.654 (0.006)	0.37 (0.009)	-0.12 (0.008)	-0.12 (0.008)
10 per cent poorest	0.053	1.89	-0.109	-1.04	2.837	1.85	-1.089	-1.44
10 per cent richest	-0.010	0.03	-0.013	0.00	0.031	0.07	-0.122	-0.03
Elasticity of the conditional expectation of hours supply, all households	0.007 (0.001)	0.11 (0.002)	-0.035 (0.002)	-0.05 (0.004)	0.078 (0.003)	0.54 (0.007)	-0.12 (0.008)	-0.12 (0.008)
10 per cent poorest	0.021	0.29	-0.017	-0.15	0.467	1.04	-1.410	-1.04
10 per cent richest	-0.030	0.03	-0.015	-0.01	0.004	0.12	-0.600	-0.06
Elasticity of the unconditional expectation of hours supply, all households	0.053 (0.002)	0.28 (0.005)	-0.116 (0.002)	-0.08 (0.005)	0.737 (0.006)	0.91 (0.11)	-0.24 (0.014)	-0.24 (0.014)
10 per cent poorest	0.075	2.23	-0.126	-1.18	3.441	3.09	-1.454	-2.23
10 per cent richest	-0.041	0.06	-0.029	-0.01	0.035	0.19	-0.181	-0.08

Note: The numbers are gathered from Aaberge et al (1999b). Standard deviations in parentheses (not reported for the income dependent elasticities).

We observe that in both countries the uncompensated labour supply elasticities decline strongly with income. For the richest households the labour supply elasticities for both genders are close to zero, and they are even negative for males in Italy. Thus, like in Dickens and Lundberg (1993) the labour supply curve is estimated to be upward sloping for households with low household income but more steep for the richest households and even backward bending in the Italian case. Wage rates for men and women typically vary in tandem over the business cycles. Hence, if labour supply

responses to economic incentives are estimated on aggregate time series it is really the net effect of changes in both these wages that is captured. The net elasticity may be defined as the own wage elasticity plus the cross wage elasticity. As seen from Table 2 the cross wage elasticities are negative and rather sizeable. Consequently the net elasticities are much lower than the own wage elasticities, in particular for poor households. Thus the decline of the elasticities with household income is somewhat weaker when judged by the net elasticities rather than by the own wage elasticities.

So far we have reported uncompensated elasticities. From a pure efficiency point of view, it may be argued that compensated (Slutsky) elasticities are more relevant. Aaberge et al (1993) report compensated elasticities for Italy in 1987, based on a family labour supply model of the type discussed above. The compensated elasticities exhibit the same variation with household income as the uncompensated do. The decline with income is *even stronger*. The own compensated participation elasticity for males is estimated to 0.052 for the 10 per cent poorest and only 0.004 for the 10 per cent richest households. The male hours-of-work elasticity is 0.176 for the poorest and 0.015 for the richest. For women, the estimated compensated participation elasticity is as high as 9.621 for the 10 per cent poorest and only 0.082 for the 10 per cent richest households. The female hours-of-work elasticity is 1.261 for the poorest and 0.014 for the richest.

3.5 Discrete Labour Supply

The evidence referred to in the previous section suggests that the literature's general preoccupation with *marginal* labour supply adjustments may be misplaced. For a large number of workers, the labour supply decision is really a discrete one. Many jobs are defined as full- or part time, with the hours of work more or less predetermined. The reasons for this may be technological (the production process requires workers to co-ordinate their activity and therefore be present simultaneously) or institutional (legislation, tariff agreements). In any case, the probability of being offered a full time job is normally much higher than the probability of being offered a job with less hours of work. The fact that in most countries taxes are not uniformly progressive adds to the need of comparing utility levels at discrete points rather than deriving labour supply from equating locally the marginal rate of substitution and the

net marginal wage rate.

The issue of having or not having a job is for many workers clearly the most pertinent labour supply decision to make. And although some hours-of-work alternatives may be available in the market, the issue is fundamentally of discrete nature. For example, the typical labour supply decisions faced by a married couple are first whether to have one or two incomes, and second, whether a second income shall be based on full- or part time work. Similar decision problems are likely to be important for single persons with low (actual or potential) market wages, relative to alternative non-market sources of income. Low-wage workers may have access to social security assistance or unemployment benefits that are relatively high compared to their expected market wage. Older workers may have access to early retirement schemes, that also involve higher replacement ratios, the lower is the income forgone. These transfers are typically lost (partly or fully) if the persons in question take up paid work. Thus, the formal tax rates for these workers are only partial; if the social security system is taken into account, low-wage workers often face very high marginal tax rates, as well as large fixed costs associated with participation[5]. This may be one of the driving forces behind the high labour supply elasticities estimated for low-wage men (Juhn et al, 1991). If the accessible take-home wage increases beyond a certain point, the labour supply jumps from zero hours to a full time job. It is also likely that low-wage jobs on average are more dull than high-wage jobs, hence the pecuniary remuneration is considered relatively more important.

Even for high-skilled prime aged workers that do tend to work full time irrespective of the tax schedules they face, an important part of the behavioural response to the income tax system is discrete in nature. For these workers, the relevant discrete decision problem is not that of labour force participation. It is the selection of a particular job. The idea that each worker faces a given wage rate determined by individual human capital only does not hold water in real economies. There is overwhelming evidence indicating that there exists industry- as well as firm specific wage differentials that are not fully attributable to differences in unobserved ability (see e.g. Gibbons

[5] If it turns out that low-wage workers in reality face higher total marginal tax rates than high-wage workers, then there may be a direct efficiency gain associated with steeper tax progression, even when the supply elasticities are the same (see. Ballard, 1988, for some illustrative calculations).

and Katz, 1992; Blackburn and Neumark, 1992). A more realistic assumption is that each worker faces an opportunity set containing a number of feasible market and non-market opportunities, each of them being characterised by a set of attributes, such as hourly wage rate, regular working time, career prospects, working conditions, type of tasks etc. The more progressive is the tax system, the less weight will be attached to the pre-tax wage rate, relative to other job attributes. Hence, the tax system may affect the types of jobs that workers choose. Results reported in Aaberge et al (1998a), based on a discrete labour supply model, confirm that there is a weak tendency for workers to select better paying jobs when the degree of tax progressivity is reduced. To the extent that pre-tax wages reflect the true value of the productive contribution in each alternative job match, higher tax progressivity thus yields a less efficient allocation of workers into jobs. However, wage differentials do not always reflect productivity, but rather the degree of market power possessed by firms or workers (unions). There is, for example, evidence that the degree of unionisation is a significant factor in explaining cross-country differences in industry wage differentials (Gittleman and Wolff, 1993).

In applied welfare analysis the tradition has been to assess the welfare effects of tax reforms based on labour supply elasticities. The fact that in most countries the budget sets are non-convex and that there are restrictions on hours offered in the market imply that the traditional approach is not sufficient in analysing the welfare effects of tax reforms. Consider for example the efficiency implications associated with a tax reform that replaces a progressive tax structure with a flat-tax system. Empirical evidence suggests that the merits of this reform (in the static sense discussed in this paper) primarily depend on the way it affects work incentives for low-skilled persons in general and married couples in particular. The majority of other workers seem to have an almost completely inelastic labour supply. Given that the labour supply choices taken by the workers of prime interest are discrete in nature, the important issue is how the alternative tax systems affect the payoff associated with a full- or alternatively a part time job relative to the payoff from non-participation. Since the workers (or households) in question typically have lower-than-average incomes, it is tempting to conclude that a flat-tax system must have even more detrimental effect on work incentives than a progressive tax system, as a flat-tax reform is likely to increase the average as well as the marginal tax rate corresponding to the income levels that are

typically observed for these workers. On the other hand, the same flat-tax reform may reduce the average as well as the marginal tax rate associated with alternative better paying (full-time) jobs that also may be available for some of these workers (if they work hard enough). If we believe that the agents in question choose between a number of discrete alternatives, the efficiency, as well as the distributive properties of alternative tax system depend on how they affect the utilities and tax revenues associated with these alternatives (Aaberge et al, 1995; 1998). If the flat-tax system tends to reduce distortions in the relative remuneration of say full participation versus no participation it improves participation-efficiency. If it increases these distortions it harms participation-efficiency.

To fix ideas consider a flat-tax reform that replaces a progressive tax system consisting of two tax brackets with one single rate (somewhere between the two rates in the existing system). To avoid inessential complications, assume that incomes other than labour income (such as unemployment benefits) are fixed net-of-taxes. Denote by Y^* the labour income level from which the higher tax rate in the existing system starts to bite, and denote by Y^{**}, the income level at which the two tax systems yields the same average tax rate. The flat-tax reform implies that the marginal tax level is reduced for $Y>Y^*$ and increased for $Y<Y^*$, while the average tax rate is reduced for $Y>Y^{**}$ and increased for $Y<Y^{**}$. For persons that choose to work anyway, and only adjust the number of hours according to the take-home wage, the flat-tax reform reduces individual dead-weight losses as long as the income is above Y^*. But how does the flat-tax reform affect participation decisions? For persons with (potential or actual) income associated with participation (the best of the discrete alternatives) above Y^{**}, participation is encouraged by the reform, while for persons with income below Y^{**}, participation is discouraged. Hence, the flat-tax reform involves smaller dead-weight losses for some workers and larger dead-weight losses for others, depending on the position of the breaking point Y^{**}. If the breaking point in the existing progressive tax system is set such that the (potential) income of a typical non-participant (or would-be non-participant) is below Y^{**}, the flat-tax reform is likely to harm, rather than enhance efficiency. If the non-participants (and those considering it) on average have lower wages than other workers, this is indeed likely to be the case. On the other hand, there may be some high-wage non-participants, particularly married women, for which progressive taxation discourages labour market participation and harms effi-

ciency. If the two spouses are taxed together, this latter case is likely to be relevant even at lower wage rates.

In short, there is no one-dimensional relationship between tax progressivity (or the degree of redistribution) and efficiency. An important challenge for policy makers is to reform features of the tax system that discourages participation. From a pure static efficiency point of view, it is probably much more to gain from removing inefficiencies in the participation decisions than from removing inefficiencies in the hours-of-work decisions.

4 Wage Determination and Equilibrium Unemployment

So far, we have built on the assumption that workers are free to choose whether or not to work. In reality, there is some involuntary unemployment. In one way or another, the reason for this is that wages don't clear the labour market. This may happen because wages have other important roles to play, such as motivating the employees (efficiency wage models), because employers or employees organise in collective wage bargaining (trade union models) or because search and matching costs furnish employers and employees with individual market power (search models). Unemployment may arise either because the aggregate wage is above the market clearing level, or because relative wages fail to reflect relative differences in demand and supply. The tax system may play a prominent role, as it affects the general wage level as well as the pre- and post tax wage distribution. Hence, it may also affect the level of unemployment. In all the models of non-competitive wage setting, it is likely that the equilibrium rate of unemployment is above its social optimum. Therefore, a complete evaluation of the efficiency loss attached to alternative tax systems must involve a consideration of how they affect wage setting. In Europe, where the level of involuntary unemployment is very high, and far above anything that can be viewed as efficient, these factors are likely to be of large quantitative importance.

Again, the literature is very much preoccupied with the representative agent framework. Such models are helpful in order to illustrate some fundamental mechanisms related to how the structure of a tax affects wage setting when the level of the tax can be taken as given. But they are not very useful for evaluating the likely wage re-

sponses arising from a tax reform that redistributes the burden of taxation between individuals. A lot of the empirical evidence presented within this area builds on equations derived from a representative agent model, but applied on data collected from a highly heterogeneous population, and it must be interpreted in that light.

4.1 The Representative Agent

Consider first how the degree of tax progressivity affects wage setting in a competitive model: Increased tax progressivity, for a given tax level, unambiguously results in a higher pre-tax wage. The reason is that the labour supply declines, as agents respond to the lower marginal take home wage by substituting towards more leisure (as the tax level is kept fixed, there is no income effect). Hence, in a competitive model, the way in which taxes affect labour supply and the way in which they affect wage setting are simply two sides of the same coin.

Non-competitive models of the labour market typically start out from a very different assumption: The labour supply is fixed. Consequently, wages are not determined directly through the interaction of demand and supply. Instead, they are set relative to unemployed workers' utility, which again depends on their current living standard and on their chances of getting a job (the rate of unemployment). It turns out that the degree of tax progressivity plays a very different role in non-competitive, as opposed to competitive, wage setting models. Indeed, within the representative agent framework, it is a common theoretical result that more progressivity reduces pre-tax wages, and thereby also reduces the equilibrium rate of unemployment. An important prerequisite underlying this unambiguous result is that a change in tax progressivity does not affect the well-being of the unemployed directly. Hence, it is assumed either that the unemployment benefit is not taxed (which is not in accordance with taxation rules in most countries), or that changes in the tax system are accompanied by adjustments in unemployment benefits (and other welfare payments), such that they are kept fixed in post-tax terms.

In a bargaining model (either collective or individual), a marginal increase in progressivity (a reduction in RIP), while keeping the tax level fixed, implies that a larger fraction of any given wage *increase* is seized by a third party (the tax collector), hence

with respect to the joint surplus of the worker(s) and the firm(s), wage increases become more expensive. Since the income effect is eliminated through the assumption of a given tax level, the bargained wage falls. In fact, tax progressivity reinforces the bargaining position of the firm relative to the worker. Consider a bilateral generalised Nash bargaining game with b representing the bargaining power of the worker (and (1-b) is the bargaining power of the firm). Then, under fairly general assumptions, it can be shown (Strand, 1999) that the bargaining outcome in a situation with progressive taxes is the same as in a situation with a flat-tax system and with the bargaining power of the worker given by $b^* = bRIP/(bRIP + (1-b))$. Hence, the effect of tax progression ($RIP<1$) is simply to weaken the bargaining position of the worker.

In a collective bargaining framework, there is an additional effect arising from the presumption that unions not only have preferences for high wages, but also for secure jobs (Hersoug, 1984). There is a trade-off between these two aims, as a higher wage typically reduces the demand for labour, and increases the risk of job losses. A more progressive tax implies that the aim of high wages has become more costly to fulfil in terms of forgone job security. This change in the relative price entails a standard substitution effect; the union substitutes the now cheaper good – job security – for the more expensive good – high wages. Hence, even in the case of a monopoly union with complete control over wage setting, more tax progressivity reduces the wage[6].

In an efficiency wage model, a marginal increase in progressivity (a decrease in RIP), ceteris paribus, makes it more costly for firms to use the wage as a tool for improving the efficiency of the employees. The typical result is, even in this case, that the optimal pre-tax wage is reduced (Hoel, 1990)[7]. As the efficiency-promoting role of the wage is played down, its market-clearing role is nurtured, which in turn reduces the equilibrium level of unemployment.

---

[6] Koskela and Vilmunen (1996) show that the negative relationship between tax progressivity and the wage level is very robust in unionised labour markets. It holds in monopoly models, in right-to-manage-models and in efficient bargaining models.

[7] Note however, that in the simple shirking model (Shapiro and Stigliz, 1984), the wage is not affected by a change in progressivity (Pissarides, 1998). The reason is that the model is discrete; workers either shirk or they don't shirk. The firms offer a take home wage that gives the premium (over the utility of being unemployed) required to deter shirking, but nothing more. Hence, only the average tax rate matters.

4.2 Heterogeneity, conflicting interests and income effects

In reality, there is of course no such thing as a representative agent. The tax reform that offers all workers the same change in residual income progression (RIP), while keeping their average tax rate fixed does not exist. Any conceivable tax reform redistributes the tax burden between individuals, at least to some extent. Some workers are made better off while others are made worse off. Hence, the substitution effect discussed in the previous subsection is accompanied by various income effects. Moreover, a tax reform that is considered to reduce the overall degree of progressivity (say by reducing the highest marginal tax rates) increases the RIP for some high wage workers, but it may very well leave it unchanged or even decrease it for others. Hence, with respect to wage setting, even the substitution effect may work in different directions for different workers. In order to evaluate the relationship between taxes and wages, there are at least two questions that need to be answered. First, how does the existence of heterogeneity interfere with the way tax progressivity affects the average wage level? And second, how does the relative distribution of taxes affect the relative distribution of pre-tax wages? To our knowledge, the literature is silent on these questions.

To illustrate the difficulties associated with these complications, assume that wages are set through a bargaining between a union and a firm. A tax reform is introduced, that increases the marginal as well as the average tax rate for the high-wage union members (reduces the RIP), while it lowers the marginal as well as the average tax rate for the low-wage union workers (increases the RIP, or leaves it unchanged). For the high-wage workers, the lower RIP provides (ceteris paribus) incentives for reducing the wage demand, while the higher average tax rate provides incentives for increasing the wage demand. The converse mechanism applies for low-wage workers. At the same time, seen from the union as a whole, it becomes more costly (in terms of job losses) to raise the wages of their high-wage members relative to the low-wage members. On the other hand, if the union pursues egalitarian objectives, a given post-tax wage distribution can be obtained with less pre-tax wage redistribution. It is very difficult to predict how these various mechanisms may interact, both with respect to the level of the average wage and with respect to the distribution of wages. The results

are likely to depend on the way workers are organised (one single union or different unions at different wage levels), as well as on the way conflicting interests are resolved within the union(s).

The wage/tax-literature is very much focused on the relationship between tax progressivity and the aggregate wage level, and how this mechanism affects the equilibrium rate of unemployment. However, the persistently high rates of unemployment in Europe have contributed to a shift in the focus of modern unemployment theory; from *aggregat*e wage setting to *relative* wage setting. This shift is based on the simple observation that unemployment in Europe is extremely unequally distributed. Workers that, if they were to be employed, would receive a fairly low wage, are much more likely to be unemployed than workers who typically receive a high wage. There are two possible explanations. First, persons with access to only low wages have weaker incentives to get a job; hence they seek work less actively, if at all. The existence of unemployment benefits and other welfare payments may, for these particular workers, erode the utility gain that otherwise would be associated with obtaining a job. Second, the egalitarian wage policy that has been pursued in most European countries may have raised the wages of the low-skilled above the firms' reservation level. For some workers, the negotiated wage may simply exceed their expected productivity and leave the workers concerned almost unemployable (Røed, 1998).

One of the most popular explanations for the current European unemployment problem runs as follows (see e.g. OECD, 1994; Krugman, 1994): There is a global trend towards increased productivity dispersion that brings about a skill-biased change in labour demand. This is the reason why we have observed a spectacular increase in *wage inequality* in the United States (see e.g. Katz and Murphy, 1992). European countries (with the exception of the UK) have typically avoided the increase in wage inequality, due to a more centralised and egalitarian wage system (e.g. in the form of minimum wages). But they have paid a price in terms of high long-term unemployment among low-skilled workers. Although this hypothesis is striking in its ability to explain simultaneously the most conspicuous labour market developments in both the United States (more inequality) and in Europe (more unemployment), it is controversial (for some counter-arguments, see e.g. Nickell and Bell, 1995; 1996; or Jackman, et. al, 1997). But if it does convey something of importance, it suggests a potential

role for tax reform in reducing unemployment. What the hypothesis implies, is that wages for some low-skilled workers are too high (including taxes) for the firms to employ them, while they are too low (net of taxes) for the workers to really want employment. A shift in the tax burden from those with low wages to those with high wages (i.e. a more progressive tax schedule) may in principle eliminate (or reduce) both these employment barriers. It may make it more attractive for low-skilled workers to work, while at the same time give the firms stronger incentives to hire them. The important question is then: Will more redistribution through the tax system really entail less redistribution through the setting of pre-tax wages in a typical European collective bargaining framework? As noted above, the answer to this question depends on a number of factors, among them the degree to which different workers are organised in different unions and the way conflicting interests within unions are resolved.

4.3 Empirical Evidence

The empirical literature about tax progressivity and wage-setting abstracts from these difficulties. The typical approach is to regress some measure of the average wage on a series of explanatory variables, among them a measure of tax progressivity. The most common approach is to take the arithmetic mean of individual RIP's or the RIP associated with the mean wage as the progressivity measure. Hence the literature is preoccupied with how (some notion of) average progressivity affects the average wage level. There is, as far as we know, not reported any empirical results about how the degree of tax progressivity affects relative wages. Table 3 gives a brief overview of the empirical evidence offered.

Table 3
Empirical evidence regarding tax progressivity and the wage level.

Authors	**Coverage**	**Type of data**	**Main results**
Malcomson and Sartor (1987)	Italy, 1968-80	Time series of average wages for industrial workers	More progressivity reduces the wage.
Lockwood and Manning (1993) and Manning (1993)	United Kingdom, 1956-87	Time series of average wages for manual workers	More progressivity reduces the wage.

Table 3
Empirical evidence regarding tax progressivity and the wage level.

Authors	Coverage	Type of data	Main results
Tyrväinen (1995)	Germany, Canada, France, Finland, Australia, USA, Sweden, Japan, UK, Italy 1972-92	Time series for average wages for "average production workers"	More progressivity reduces the wage in Canada, Japan, Italy and Finland. No significant effect in the other countries.
Wulfsberg (1996)	Norway 1972-91	Panel data for the average production worker wage in 7000 firms	More progressivity reduces the wage.
Bowitz and Cappelen (1997)	Norway 1968-93	Time series of average wages and tax rates for manual workers	More progressivity reduces the wage.
Dyrstad and Lund (1997)	Norway 1979-88	Panel data for average manual wages in 333 municipalities.	More progressivity reduces the wage.
Holmlund and Kolm (1995)	Sweden 1975-92	Time series of average earnings for 5 different income groups (deciles).	More progressivity reduces the wage. It is reduced most for those with highest incomes.
Holmlund and Kolm (1995)	Sweden 1989 and 1992	Microdata for individuals.	More progressivity reduces the wage.
Lockwood et. al. (1998)	Denmark 1970-92	Time series of average earnings for 9 different income groups.	More progressivity reduces the wage for low- and middle income workers, while it increases the wage for those with very high wages.
Hansen et. al. (1996)	Denmark 1970-92	Time series for low-skilled male workers and high-wage white-collar workers.	More progressivity reduces the wage for low-skilled male workers, while it increases the wage for high-wage white-collar workers.

Even though these various pieces of evidence are quite unanimous in terms of identifying a negative relationship between the degree of tax progressivity and the wage rate, there are potentially some serious problems, both with respect to the econometric techniques and with respect to economic interpretation. The most important problem is that of aggregation. The tax system is highly non-linear, and it is not at all clear how an average measure of progressivity should be calculated. The Danish results referred to in Table 3 (Lockwood et. al., 1995; Hansen, et. al., 1996) indicate that there may be very different mechanisms governing the wage setting process at different points in the wage distribution. One interpretation of the findings in Denmark is that

most workers do get their wages determined through some form of non-competitive wage setting, most likely a collective bargaining process. In the upper tale of the wage distribution (and to some extent also in the lower tale), however, wages appear to be generated in a more competitive environment, hence the labour supply effects dominate wage setting for these workers[8].

Econometric problems arise for three additional reasons. First, for the studies based on aggregate time series, it is difficult to isolate effects of changes in tax progressivity on unemployment and wages during a period of strongly increasing unemployment rates. Second, all the studies confront a serious simultaneity problem: Not only does the wage level depend on the degree of tax progressivity, the degree of tax progressivity also depends on the wage level. Although most of the studies listed above try to sort out that problem with the aid of instrumental variables, it is not unlikely that some simultaneity bias remains. And finally, unions, particularly in the Scandinavian countries, often declare that they pursue egalitarian wage policy. As there is a positive correlation between progressivity (RIP) and the wage level, it may be difficult to sort out the pure progressivity effect from the egalitarian-wage-policy effect.

5 Conclusion

The trade-off between equality and efficiency is not inevitable. The textbook unequivocal dead-weight loss associated with tax progressivity disappears as soon as the unique representative agent is replaced by heterogeneous agents or as soon as the perfect market economy is replaced by more realistic second-best alternatives.

A key element in the assessment of social efficiency is how a tax system affects participation decisions. In particular, we have pointed out that if a progressive tax system is devised in order to reduce the tax rate at income levels typically available to the low-skilled, it may encourage participation for these workers much more than it discourages hours-of-work for other workers. We believe that the marginalist approach commands too much of the way public economists, as well as policy makers, think about taxation. A tax reform proposal aiming at improved social efficiency should not

[8] Note however, that the results from Sweden (Holmlund and Kolm, 1995) tell a different story.

be judged by the degree to which it removes work-disincentives for prime aged males (who work full time anyway). It is quantitatively much more important how it affects the work-incentives for low-skilled and older workers in general, lone mothers and married women in particular (who perhaps choose to stay out of the labour force). Unfortunately, many of the tax reforms enacted during the last decades have had the wrong focus. By reducing marginal taxes primarily for the rich, they have made the rich richer, probably without delivering much in terms of improved social efficiency.

While a progressive tax may have only small effects on *marginal* labour supply decisions, it may have large effects on marginal wage setting. We have shown that progressive taxation (for a given tax level) may entail moderate wage setting in non-competitive labour markets, and hence reduce the equilibrium level of unemployment.

With reference to currently prevailing tax systems, it is probably possible to enhance efficiency, both through reforms that reduce the degree of redistribution and through reforms that increase it. The big efficiency-issue is not the degree of redistribution. But it is the way in which it is obtained.

References

Atkinson, A. B. and Stiglitz, J. E. (1980) *Lectures on Public Economics*. McGraw-Hill.

Arrufat, J. L. and Zabalza, A. (1986) Female Labour Supply with Taxation, Random Preferences, and Optimization Errors. *Econometrica*, Vol. 54, 47-63.

Ballard, C. L. (1988) The Marginal Efficiency Cost of Redistribution. *The American Economic Review*, Vol. 78, 1019-1033.

Blackburn, M. K. and Neumark, D. (1992) Unobserved Ability, Efficiency Wages and Interindustry Wage Differentials. *Quarterly Journal of Ecnomics*, Vol. 107, No. 4, 1421-1436.

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

Blundell, R. (1996) Labour Supply and Taxation. In Devereux, M. (ed.): *The Economics of Tax Policy*, Oxford University Press.

Blundell, R. (1997) The Impact of Taxation on Labour Force Participation and Labour Supply. *The OECD Jobs Study*, Working Paper Series No. 8.

Borjas, G. and Heckman, J. J. (1978) Labour Supply Estimates for Public Policy Evaluation, in Industrial Relations Institute. *Proceedings of the Thirty-First Annual Meeting*, Madison, WI, 320-331

Bowitz, E. and Cappelen, Å. (1997) Incomes Policies and the Norwegian Economy 1973-93. *Discussion Papers No. 192*, Statistics Norway.

Browning, E. K. and Johnson, W. R. (1984) The Trade-Off between Equality and Efficiency. *Journal of Political Economy*, Vol. 92, 175-203.

Burtless, G. and Hausman, J. A. (1978) The Effect of Taxation on Labour Supply: Evaluating the Gary NIT Experiment. *Journal of Political Economy*, Vol. 86, 1103-1130.

Cogan, J. (1981) Fixed Costs and Labour Supply. *Econometrica*, 49, 945-964.

Creedy, J. (1996) *Fiscal Policy and Social Welfare. An Analysis of Alternative Tax and Transfer Systems*.

Dagsvik, J.K., Laisney, F., Strøm, S. and Østervold, J. (1988) Female Labour Supply and the Tax Benefit System in France. *Annales D'Economie et de Statistique*,

Vol. 11, 5-40.

Devanzo, J., Detray, D., and Greenberg, D (1973) *Estimating Labour Supply Response: A Sensitivity Analysis*, Rand report R-1372-OEO, Santa Monica, The Rand Corporation.

Dickens, W. and Lundberg, S (1993) Hours Restriction and Labor Supply. *International Economic Review*, Vol. 34, 169-191.

Dyrstad, J. M. and Lund, M. (1997) Income Taxes and Wages: Empirical Evidence from Norway. Unpublished manuscript, Norwegian University of Science and Technology.

Ebert, U. (1992) Global Tax Progressivity. *Public Finance Quarterly*, Vol. 20, 77-92.

Gibbons, R. and Katz, L. (1992) Does Unmeasured Ability Explain Inter-Industry Wage Differentials? *Review of Economic Studies*, Vol. 59, 515-535

Gittleman, M, and Wolff, E. (1993) International Comparisons of Inter-industry Wage Differentials. *Review of Income and Wealth*, Vol. 39, No. 3, 295-312.

Hammond, P. (1990) Theoretical Progress in Public Economics. A Provocative Assessment, *Oxford Economic Papers*, Vol. 42, 6-33.

Hansen, C. T., Pedersen, L. H. and Sløk, T. (1996) Danske resultater om sammenhængen mellem marginalskat og løn (Danish results about the relationship between the marginal tax rate and the wage). *Nationaløkonomisk Tidsskrift*, Vol. 134, 153-174.

Hausman, J. A. (1979) The Econometrics of Labor Supply on Convex Sets. *Economic Letters*, ***3,*** 171-174.

Hausman, J. A. (1980) The Effects of Wages, Taxes, Fixed Costs on Women's Labor Force Participation. *Journal of Public Economics*, Vol. 14, 161-192.

Hausman, J. A. (1981) *Labor Supply*. In Aaron H. and Pechman J., (eds): *How Taxes Affect Behavior*, Washington, D.C. Brookings Institution, 27-83.

Hausman, J. A. (1985a) The Econometrics of Non-Linear Budget Sets. *Econometrica*, Vol. 53, 1255-1282.

Hausman, J. (1985b) *Taxes and Labour Supply*. In A. Auerbach and M. Feldstein (eds) *Handbook of Public Economics*, Vol. 1, Amsterdam, North-Holland, 213-263.

Hausman, J., and Ruud, P. (1984) Family Labour Supply with Taxes. *American Eco-*

nomic Review, Vol. 74, 242-253.

Heckman, J. J. (1978) A Partial Survey of Recent Research on the Labour Supply of Women, *American Economic Review*, Vol. 68, 200-207.

Hersoug, T. (1984) Union Wage Responses to Tax Changes. *Oxford Economic Papers*, Vol. 37, 37-51.

Hicks, J. (1939) *Value and Capital*. Oxford: Clarendon Press.

Hoel, M. (1990) Efficiency Wages and Income Taxes. *Journal of Economics*, Vol. 51, 89-99.

Holmlund, B. and Kolm. A. S. (1995) Progressive Taxation, Wage Setting, and Unemployment: Theory and Swedish Evidence. *Swedish Economic Policy Review*, Vol. 2, 423-460.

Ilmakunnaas, S. and Pudney, S. (1990) A model of female labour supply in the presents of hours restrictions. *Journal of Public Economics*, Vol. 41, 183-210.

Jackman, R., Layard, R., Manacorda, M. And Petrongolo, B. (1997) European Versus US Unemployment: Different Responses to Increased Demand for Skill? Centre for Economic Performance, Discussion Paper No. 349.

Jakobsson, U. (1976) On the Measurement of the degree of Progression. Journal of Public Economics, Vol. 5, 161-168.

Juhn, C., Murphy, K. M. and Topel, R. (1991) Why Has the Natural Rate of Unemployment Increased over Time? *Brookings Papers on Economic Activity*, 2, 75-126.

Katz, L. F. and Murphy, K. M. (1992) Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *Quarterly Journal of Economics*, Vol. 107, 36-78.

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

King, M. (1987) *The Empirical Analysis of Tax Reforms*. In T. Bewley (ed): Advances in Econometrics, Cambridge University Press, Cambridge.

Koskela, E. and Vilmunen, J. (1996) Tax Progression is Good for Employment in Popular Models of Trade Union Behaviour. *Labour Economics*. Vol. 3, 65-80.

Krugman, P. (1994) Past and Prospective Causes of High Unemployment. Lecture held at a symposium i Jackson Hole, Wyoming, august 1994.

Leibfritz, W., Thornton, J. and Bibbee, A. (1997) *Taxation and Economic Performance*. OECD, Economics Department, Working paper 176.

Lockwood, B. and Manning, A. (1993) Wage Setting and the Tax System. Theory and Evidence for the United Kingdom. *Journal of Public Economics*, Vol. 52, 1-29.

Lockwood, B. Sløk, T. and Tranæs, T. (1998) Progressive taxation and wage setting: evidence for Denmark. Unpublished manuscript, Copenhagen Business School.

Malcomson, J. M. and Sartor (1987) Tax Push Inflation in a Unionized Labour Market. *European Economic Review*, Vol. 31, 1581-1596

Manning, A. (1993) Wage Bargaining and the Phillips Curve: The Identification and Specification of Aggregate Wage Equations. *The Economic Journal*, Vol. 103, 98-118.

Mirrlees, J. A. (1971) An exploration in the theory of optimum income taxation. *Review of Economic Studies*, Vol. 38, 175-208.

Moffitt, R. A. and Wilhelm, M. (1998) *Taxation and the Labour Supply Decisions of the Affluent*. NBER Working Paper 6621.

Mroz, T. (1987) The Sensitivity of an Empirical Model of Married Women's Hours of Work to Economic and Statistical Assumption. *Econometrica*, Vol. 55, 765-800.

Musgrave, R.A and Thin, T. (1948) Income Tax Progression, 1929-48. *Journal of Political Economy*, Vol. 56, 498-514.

Nickell S. and Bell, B. (1995) The Collapse in Demand for the Unskilled and Unemployment Across the OECD. *Oxford Review of Economic Policy*, Vol. 11, 40-62.

Nickell S. and Bell, B. (1996) Distribution of Wages and Unemployment in OECD Countries. *American Economic Review*, PAP, Vol. 86, 302-308.

OECD (1994) *The Jobs Study*. Paris

OECD (1995) *The OECD Jobs Study. Taxation, Employment and Unemployment*. Paris

Pissarides, C. A. The impact of employment tax cuts on unemployment and wages; The role of unemployment benefits and tax structure. *European Economic Review*, Vol. 42, 155-183.

Reynolds, M. and Smolensky, E. (1977) Post-Fisc distributions of income in 1950, 1961 and 1970. *Public Finance Quarterly*, Vol. 5, 419-438.

Røed, K. (1998) Egalitarian Wage Policies and Long Term Unemployment. *The Scandinavian Journal of Economics*, Vol. 100, No. 3, 611-625.

Rønsen, M. and Strøm, S. (1991) *Enslige forsørgeres tilpasning mellom trygd og arbeid*. In Aksel Hatland (ed) Trygd som fortjent, Ad Notam.

Shapiro, C. and Stiglitz, J. (1984) Equilibrium Unemployment as a Discipline Device. *American Economic Review*, Vol. 74, 433-444.

Stiglitz, J. (1986) *Economics of the Public Sector* (second edition). W.W. Norton.

Strand, J. (1999) The effects of progressive taxation under decentralized bargaining and heterogeneous labour. Unpublished working paper. Department of Economics, University of Oslo.

Sørensen, P. B. (1997) Public finance solutions to the European unemployment problem? *Economic Policy*, October 1997, 221-264.

Triest, R. K. (1994) The Efficiency Cost of Increased Progressivity. In Slemrod, J. B. (ed.): *Tax Progressivity and Income Inequality*. Cambridge University Press.

Tyrväinen, T. (1995) Real Wage Resistance and Unemployment: Multivariate Analysis of Cointegrating Relations in 10 OECD Countries. *The OECD Jobs Study Working Paper Series*, No. 10.

Walker, I. (1990) The Effects of Income Support Measures on the Labour Market Behaviour of Lone Mothers. *Fiscal Studies*, Vol. 11, No. 2, 55-75.

Wulfsberg, F. (1996) Do Progressive Taxes Reduce Wage Pressure? Chapter 4 of: Panel Data Evidence on Wage Setting and Labour Demand from Norwegian Manufacturing Establishments. PhD thesis, University of Oslo.

Aaberge, R., Colombino, U. and Strøm, S. (1993) Labor Supply in Italy. *Discussion Papers 92*, Statistics Norway.

Aaberge, R., Dagsvik, J.K. and Strøm, S. (1995) Labour Supply Responses and Welfare Effects of Tax Reforms. Scandinavian Journal of Economics, Vol. 97, No. 4, 635-659.

Aaberge, R., Colombino, U. And Strøm, S. (1999a) Labour Supply in Italy. *Journal of Applied Econometrics*, forthcoming.

Aaberge, R., Colombino, U. and Strøm, S. (1999b) Labor Supply Responses and Welfare Effects from Replacing Current Tax Rules by a Flat Tax: Empirical Evidence from Italy, Norway and Sweden, forthcoming in *Journal of Population Economics.*