

The Public Economics of Redistribution and the Welfare State

Agnar SANDMO*

Abstract This article gives a broad survey of the arguments why governments should engage in large-scale redistribution, as they do, e.g., in the Scandinavian countries. Official data for redistribution via government budgets tend to overstate its magnitude, since much of the aggregate volume is due to income smoothing, which is really a substitute for private saving and insurance. The arguments in support of socialized saving and insurance must basically be about market failure, while the amount of net interpersonal redistribution must be based on a calculation of the trade-off between efficiency and equity. The optimal degree of income equalization depends on the efficiency cost of the tax system.

1. Introduction

How far should we go in the equalization of incomes? This is perhaps the fundamental question that faces policy makers with a commitment to egalitarian values. A fervent egalitarian would perhaps argue that the right thing would be to proceed all the way to complete equality. On second thought, however, most people would probably admit that such a policy would have strongly adverse effects on economic incentives, which should act as a brake on redistribution policy. How far we should go then becomes a question partly of ethics, partly of economics. I would like to emphasize that I certainly do not claim—neither on my own part nor on the part of the economics profession—to know the precise answer to the question. But I note as a fact that there is a lot of popular and political support for the redistribution schemes of the welfare state. This is revealed in many sources, as in the programs of the political parties and in public opinion polls. Taking this support for a fact, my analysis in the following consists in trying to think systematically about redistribution policy and its relationship to other objectives of public policy. It does not imply a strong commitment to any specific ethical or philosophical view, although I shall occasionally use formulations derived from such views to illustrate my general points.

Starting from this basis. I would like to consider the question of how far it is

* Professor of Economics, Department of Economics, Norwegian School of Economics and Business Administration, Norway.

I am grateful to Mats Persson for his comments on an earlier version of this article.

reasonable to go in the direction of income equality. To do this, I try and assemble bits and pieces of economic theory that together may lead to a hopefully coherent view of the role of redistribution and provision of social insurance in the welfare state.

What that role is, has been the subject of a lively debate over the years, so that it is by no means a trivial matter to arrive at a balanced view. The title of a famous book by Arthur Okun (1975), *Equality or Efficiency: The Big Trade-off*, probably summarizes the widely held conviction that more equality can only be bought at the price of less efficiency, roughly in the form of a reduction of GDP or GDP growth. However, not all economists agree with this; several writers, e.g. Nicholas Barr (1987), have emphasized that the welfare state can also be interpreted as a device to overcome some basic market failures in the private sector. From this perspective, therefore, there may be grounds for doubting whether every aspect of welfare state policy really involves any trade-off.

2. How Much Redistribution Is There?

I am not going to dispute the view that the welfare state policies of the Scandinavian countries (which are the ones that I know best) really achieve substantial redistribution. If you use annual data to compute Lorenz curves for factor incomes on the one hand and for disposable incomes on the other, they typically look like the picture shown in Fig. 1. Not only does the Lorenz curve for factor incomes lie below the curve showing the inequality in disposable income; especially for Sweden it lies in fact substantially below, as demonstrated by Fritzell (1991). Fritzell shows in particular that the Gini inequality index in 1980 was only half as large for disposable incomes as it was for factor incomes, and he refers to this difference as the redistribution effect. I want to begin my analysis by discussing whether that is really a good measure of the actual redistribution that the welfare state achieves. This is not meant to detract from the value of studies like that of Fritzell and others. Empirical research is necessarily constrained by which data are available, and I really want to argue that we should also be interested in data that—at least in large measure—are unobservable.

Let us imagine that we start from data about the distribution of labour incomes. This, of course, shows great inequality. Not only is there considerable inequality in the distribution of income among those who are actually in the labour force, but even more importantly, there is a large number of people (approximately half the population) who do not work at all. This group includes children and young people who are being educated, the old, the sick and the disabled. It also includes a number of people who have voluntarily chosen not to work, such as mothers with young children.




Figure 1 Lorenz curves

The reason why a large number of the people who do not have any labour incomes, still have a reasonable standard of living, is not only that the welfare state redistributes income in their favour. Some individuals live on income from capital. Another (related) reason is that consumption is not only determined by current income, but also by past and future income. A study of the distribution of lifetime income would almost by definition show less inequality than that of annual income; for an empirical analysis see Lindbeck (1983). A further reason is that some people are taken care of through the voluntary support of their families and friends.

At a more general level it seems clear that if we are concerned with inequality, what really matters is not the distribution of income per se, but the distribution of the *standard of living* between persons. That is not an easy concept to make empirically operational. It clearly depends on the level of consumption of private goods, but also on the supply of public goods (public safety, the physical and cultural environment) and publicly provided private goods, such as health and education. It also depends on the amount of voluntary transfers received from others, above all through the family. In the following I shall concentrate on private goods consumption, so let me first say a few words about the components of the standard of living that are not captured by data on private consumption.

It is clear that the public provision of certain goods and services makes the distribution of the standard of living less dependent on income. Although there may

in fact be a positive relationship between the access to and use of certain public services on the one hand and income on the other, it still seems safe to assume that the value of such services is more equally distributed than labour incomes. From this point of view, since Sweden—the prime example of a highly developed welfare state—has more public provision of both public and private goods than most other countries, one might perhaps expect that the equalization of incomes would be of less importance in Sweden. In fact, to the extent that inequality measures are really internationally comparable, there is more redistribution of income in Sweden than in the rest of the OECD area. I leave this as a puzzle to be reflected upon while emphasizing the general point that the deeper study of inequality should also take account of the redistributive impact of public expenditure.

The distribution of factor incomes is in principle observable. Now let us imagine that we were able to observe what the distribution of private income and consumption would have been, had the public sector *not* engaged in redistribution. This is nonobservable, but we can make some pretty good guesses at what it would look like. In terms of financial resources, the old people without labour incomes would live off their savings. The young people at universities would live on private loans. Disabled people would live on income provided through disability insurance. Within all of these groups we would in addition observe a considerable amount of private redistribution, partly through organized charity, partly through intrafamily transfers. Thus, in this hypothetical situation the old and disabled would to a larger extent than now live with their families, and students' expenses would more likely be paid by their parents.¹ To summarize, the market's allocation of consumption, i.e., the private component of the standard of living, would be considerably less unequal than the distribution of labour income.

Note that both private and public redistribution are of two kinds. When people save for their old age or insure against work disability, they engage in *intrapersonal redistribution or income smoothing*, i.e., they transfer income between periods and states of nature² in their own lives. The transfer of income from rich to poor could on the other hand be characterized as *interpersonal redistribution*. The latter kind is what we usually think of when we talk about

¹ In judging the realism of this scenario, one must of course keep in mind that the working part of the population would have much higher disposable incomes, since the government would no longer need tax revenue for redistribution.

² I use the concept of the state of nature in the sense made familiar by the work on general equilibrium theory by Arrow (1953) and Debreu (1959). To illustrate, imagine that in some future period, my standard of living depends crucially on whether or not I am able to work. This defines two states of nature as "healthy" and "disabled". The optimum amount to save for the future depends on which state actually occurs. In this situation a rational course of action for me would be to channel some of my present saving into a form where the return is contingent on the state "disabled" occurring, i.e., a disability insurance policy.

redistribution, but it is important to be aware that much of welfare state redistribution is really of the first kind. In the case of old age pensions, one could compute the actuarial value that would correspond to the premium paid over the years spent in the labour force and compare that with the actual value of the benefits. Particularly in the universal systems that we have in the Scandinavian countries, a comparison of the differences so obtained would show the extent of redistribution of the old age pension system, but it would also show that the actuarial part of the pension is very considerable.³

In order to correctly assess the redistributive effect of the welfare state, our standard of comparison should not be the distribution of factor incomes that we actually observe, but the distribution of consumption if the state had not engaged in redistribution. Even if we cannot observe this, we should note that this in all probability would lead us to a much lower estimate of the effect. The point is that public redistribution tends to crowd out both income smoothing and voluntary transfers. While the latter will be discussed below, a stylized example of the former is shown in Fig. 2. This shows the case of an individual who lives and consumes in two periods, working in the first and being retired in the second. His first-period income is Y_1 , and in the absence of any government pension schemes he saves an amount which enables him to achieve his desired lifetime consumption profile, $C^*=(C_1^*, C_2^*)$. Now suppose that the government introduces a pension scheme under which the consumer pays a certain contribution in the first period in return for a pension in the second period. If the pension scheme is actuarially fair, this scheme moves the initial income point of the consumer—in terms of disposable income—to (say) Y' . The point to note is that the pension in this case does not change the present value of the consumer's lifetime income. Therefore, his preferred consumption profile remains the same. What changes is his private saving, which is reduced by the exact amount of the public pension premium. (If the amount of the premium had been sufficiently high, the consumer would have responded by personal dissaving, so that he would actually have borrowed against his future pension. This raises some questions about the nature of capital markets that I do not go into here.)

Suppose now that we were able simultaneously to observe two persons who were identical in terms of preferences and lifetime private incomes, except that we observed one of them in the first and the other in the second period of his

³ Mats Persson (1998) has calculated that for the Swedish working population as a whole, about 75 % of the pension charge must be considered a tax while the remaining 25 % is an insurance premium. However, these numbers do not necessarily reflect the redistributive impact of the pension system, since they refer to the *marginal* effect of the pension charge. In order to assess the balance between intra- and interpersonal redistribution in the system, one would have to study the differences between the actuarial value of the *total* benefits and contributions.




Figure 2 A public pension crowds out private saving

life. A study of labour market inequality would show this to be very large in terms of wage income, since one person would have all of the income. The distribution of consumption would show much less inequality, since the consumption of the old person would now be equal to his pension, while the income of the young would be reduced by the amount of his contribution; in fact, as the diagram has been drawn, the individuals have a strong preference for consumption equality over their own lifetime, so that consumption is virtually the same. In this case, therefore, the pension scheme would not have changed the real amount of inequality, since the distribution of consumption among the two groups is exactly the same as in its absence.

I conclude by observing that data about the amount of redistribution must be interpreted with a good deal of care. There is no doubt that the welfare state carries out much redistribution, but also that much of it is what we have called intrapersonal redistribution, so the interpersonal redistribution of income is less than what might appear from a superficial inspection of the statistics. I now turn to a consideration of the arguments for the two kinds of redistribution.

3. Government Interference in Intrapersonal Redistribution

Why would the government engage in intrapersonal redistribution (income

smoothing)? Figure 2 showed a case where there was no point in this, since the distribution of the standard of living was independent of government policy. The argument could even be extended to cover some insurance markets where similar crowding out arguments would apply. However, it should be kept in mind that the reasoning was based on some very strong assumptions that need to be examined carefully.

The basic assumption is that markets are perfect. There must exist a perfect capital market for ordinary borrowing and lending, as well as a complete set of insurance markets in which individuals may obtain any desired degree of coverage against events that may affect their standard of living. This would include such events as sickness, disability and unemployment. The model of general competitive equilibrium developed by Arrow (1953) and Debreu (1959) has the property that a competitive equilibrium is a Pareto optimum, so that a rational policy of redistribution would seem to be limited to the interpersonal kind.

However, the Arrow-Debreu model is a very special case. Let us consider the assumption of a complete set of insurance markets. In the more recent literature much attention has been paid to problems that arise out of the existence of asymmetric information, viz. moral hazard and adverse selection. Moral hazard exists when the behaviour of the insured person is influenced by the insurance. The theoretical hypothesis is that automobile insurance makes car owners drive less carefully, unemployment insurance makes people less keen to hold on to their jobs, and sickness insurance makes people less careful in the choice of their private lifestyles. Adverse selection is the case made famous by Akerlof (1970). If individuals differ with respect to their risk exposure, but the insurance company only knows the average risk in the population or the risks characteristics of broad groups, then the bad risks will drive out the good risks, possibly resulting in the insurance market for the risk in question vanishing altogether. Examples are easy to find: Insurance against unemployment would be especially attractive to those with a high risk of being laid off, while insurance against becoming a lone parent would attract people with a lifestyle that involved a particularly high risk that this event would occur. These reasons for failure of private insurance markets have been extensively discussed and are by now well understood. It is also generally accepted that social insurance—in the sense of mandatory universal membership—can overcome the adverse selection problem, but not the problem of moral hazard. A major part of the reason for market failure carries over to the case where insurance has been socialized, and the incentive problems associated with extensive social insurance coverage constitute the potentially most significant cost of the system.

As has been pointed out by Atkinson (1991), it is difficult to interpret the growth of the social security system without considering the development of the

labour markets in modern industrial societies. As a result of the development of modern industrial societies, labour has become ever more specialized, often as the result of large investments in human capital through specialized education or on-the-job training. These are investments that cannot easily be diversified through private markets, and they are also—both because of moral hazard and adverse selection problems—difficult to ensure against on a voluntary basis. The social security system provides protection against these risks, and the progressive income tax may also be seen in part as an insurance device, by which the state shares in the risk of income fluctuations; see Varian (1980).

The market failure justifications for social insurance have been extensively discussed in the literature. However, there are also other arguments, which should perhaps be given more attention than they have so far received in the welfare economics of social security and certainly in positive studies of the actual development of the welfare state. Let me briefly mention three such arguments.

The first of these is the free rider problem. Someone who realizes that he lives in a society which is committed to certain standards of equality and protection against poverty, might be tempted, if all saving and insurance were private and voluntary, to refrain from providing for himself in the expectation that society would take care of him anyway by selective social assistance programs. These programs would have to be financed by general tax revenue, but the benefits would to a large extent accrue to the free riders. Mandatory membership in the social insurance system does not completely solve this problem, since it would still be possible for someone with inadequate saving and insurance coverage to exploit the welfare system. But it no doubt relieves it considerably.

The second argument is a truly classical one, being associated with names like Böhm-Bawerk and Pigou. These and other authors argued that individuals were likely to underestimate their own future needs, discounting future utilities at a subjective discount rate that was too high; this led to a rate of saving that was too low in the sense that later in life they would come to wish that they had saved more.⁴ A public pension would—according to this argument—force people to save more that they would have done voluntarily. There is little doubt that this

⁴ Whether the rate of saving would also be too low from a social point of view is a complex issue precisely because of the non-constancy of individual preferences. - A modern version of the argument stems from the article by Strotz (1955–56), who introduced the problem of *time inconsistency*. A rational individual who plans his future time profile of consumption at time t , will—unless his intertemporal preference ordering is of a very special form—wish to deviate from the plan at time $t+1$. Realizing this, he may wish to precommit himself, e.g. in the form of a mandatory pension scheme. His analysis led Strotz to question the meaning of the principle of consumer sovereignty for intertemporal resource allocation; “the individual over time is an infinity of individuals”. See also the account of the related concept of “Allais optimality” in Malinvaud (1972, ch. 10).

argument historically has been very important for the development of the pension part of the social security system. Although the analysis in Fig. 2 casts some doubt on the effectiveness of the pension in increasing future consumption, alternative assumptions about the operation of capital markets might lead to a more positive view of this. A related line of reasoning suggests that the government might not always accept people's subjective probabilities of events like unemployment and disability. If people tend to take a too optimistic view of the occurrence of such events, one could build a case for forcing them to have a better insurance coverage; see Lundholm (1998).

The third argument concerns transactions costs. In our formal theories, building a savings and insurance portfolio is costless. But we all know that building a savings portfolio which yields both a desirable time profile of consumption and a satisfactory risk coverage is a difficult task. It may be made easier by securing expert advice, but this does not come for free. Also, the return on the portfolio depends on one's own effort and ability—in fact, just as labour income does. Mandatory saving and insurance can therefore be regarded as an attempt to economize on transactions costs, exploiting economies of scale. In this, there is also an element of redistribution, since it makes the benefits less dependent on ability and income—just as progressive taxation of wage income makes disposable income less dependent on an individual's productivity.

4. Interpersonal Redistribution

So far we have mainly been concerned with governmental interference with decisions concerning intrapersonal redistribution through saving and insurance decisions. From this point of view the aim of welfare state policies is to make time and risk markets function more efficiently. The welfare state can therefore be seen as attempt to improve upon the market system in efficiency terms.

However, this is clearly not the whole story. The welfare state also involves, both through social security and the tax system, a substantial volume of interpersonal redistribution. E.g., a more realistic picture of the pension system is the one shown in Fig. 3. Here the high-income person gets a benefit which is less than actuarially fair, while the low-income person gets more than his actuarial expectation. There is accordingly a redistribution from high-income to low-income individuals through the social security system, and this comes in addition to what is achieved through the progressive tax system.

In our account of public interference with intrapersonal redistribution the fundamental justification for public policies lay in some basic assumption about the nature of preferences, particularly with regard to risk aversion, in combination

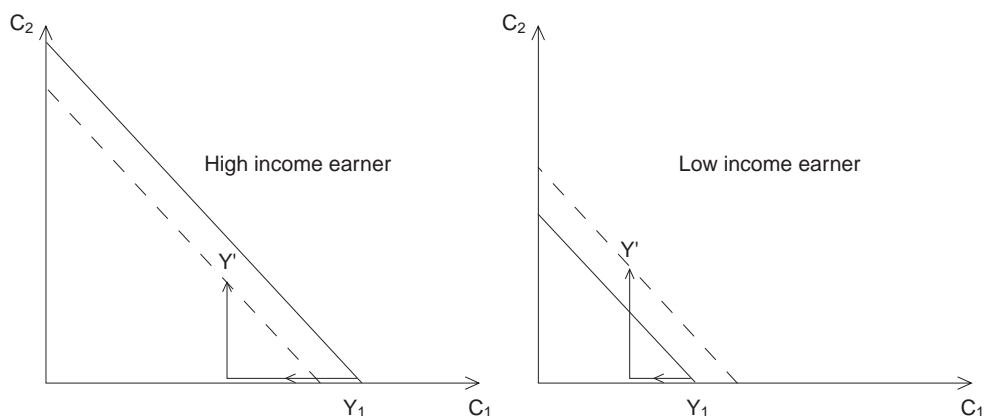


Figure 3 A redistributive pension scheme

with an assumption about incomplete markets. The existence of gambling notwithstanding, it is fair to say that the prevalence of risk aversion is well documented through actual behaviour, so that helping to diversify risk more efficiently can be interpreted as helping to satisfy people's preferences. In the case of interpersonal redistribution it is less clear that one can appeal to such revealed preference arguments. The amount of voluntary income transfers that we actually observe is clearly not insignificant, but we do not know what that amount would have been, had not the transfers been made one of the main responsibilities of the state. Do the transfers generated by the welfare state crowd out private philanthropy?

Consider a model where individuals have preferences for transferring income to others. Assume first that person A is altruistic in the sense that his utility depends both on his own consumption and on that of B. Assume that A is rich and B is poor, and that A decides to give some of his income to B. Now the state decides on a policy that increases the income of B by taxing A. If the tax paid by A equals the transfer received by B, the model implies that A will reduce the amount of his voluntary giving by exactly the amount of the transfer. The welfare state does indeed crowd out private philanthropy and achieves in fact no redistribution at all.⁵ This confronts us with a puzzle: If this analysis is even just an approximation to reality, why is there public redistribution? The puzzle may be resolved by reformulating the model to look at the problem from the point of view of a single A as before, but assuming that there is a large number of Bs, so

⁵ It is doubtful whether this type of model can satisfactorily explain the observable facts about charitable giving. For criticism and further discussion see Sugden (1982) and Andreoni (1990). Both these authors point out that a gain in realism may be obtained by assuming that donors also derive utility from the act of giving.

that we try to explain the behaviour of a single rich donor who is confronted with many poor recipients. If A's donation is divided equally among the Bs, it makes little difference to any single B. If, at the same time, A realizes that there are many like him, the problem assumes a public good character, whereby each A has incentives that imply that too little is offered as private donations.⁶ The problem might be overcome by a social contract whereby the rich As commit themselves to contribute jointly to the welfare of the poor. This contract has an obvious interpretation as the welfare state.

Another justification and explanation for redistribution policy, represented most prominently by the work of Vickrey (1945), Harsanyi (1955) and Rawls (1971), consists in deriving the desire for income equalization from the theory of risk aversion. The basic idea of Vickrey and Harsanyi is the following: Imagine that you were asked to choose between participating in a number of lotteries, where each lottery gives you an equal chance of becoming any one person in a particular country. How would you evaluate the utility of each lottery? Appealing to the axioms of the von Neumann-Morgenstern expected utility theorem, they show that the value to you of each of the lotteries would be given by the expected utility associated with it, or, since the probability of becoming any one person is the same for all, by the sum of utilities, which is the utilitarian social welfare function. If you are risk averse, you would also, in the choice between two lotteries/countries whose average income level is the same, choose the one with the least unequal distribution of income. In other words, risk aversion implies inequality aversion. Rawls' argument is similar, although his notion of risk aversion leads to the welfare criterion of the wellbeing of the least advantaged person in each society. Thus, the general argument is that risk aversion implies aversion against inequality, provided that you evaluate society from behind a veil of ignorance concerning your own position in it. There is a difficulty in making this idea operational, since as a rule you do in fact know your own position in society. Perhaps you could come close to the Rawlsian "original position" by thinking about the organization of society from the point of view of your children and children's children and so on, some of whom are not yet born and whose preferences and abilities cannot therefore be known, but whom you still care about. From this point of view, it is quite reasonable to assume that you would like to see a society with some system of redistribution from the more to the less fortunate members of society.

Given that such a social preference for equality could be established, how far would it be rational to go? The general answer to this question is that it depends firstly on the precise formulation of the preference for equality (the social welfare

⁶ This is a version of Sen's "isolation paradox"; see Sen (1967).

function), and secondly on the policy instruments that are available for redistribution. Let us suppose that the policy objective can be represented as the maximization of the utilitarian sum of utilities. Suppose also that the government can redistribute income through lump sum transfers, i.e. taxes and subsidies that are related directly to personal characteristics and that do not affect any relative prices as individuals see them. As a utilitarian you would also like to assume that the marginal utility of income is decreasing, reflecting the judgement that a dollar of extra income is worth less to a rich than to a poor person. The optimal amount of redistribution will then be given by the condition that the marginal utility of income is the same for all. Let us think about this in the simplified setting in which there are two homogeneous groups of individuals, the workers and the disabled, and that disability pensions can be financed through a lump sum tax on workers. Let m be the marginal utility of income, and let superscripts 1 and 2 refer to the workers and the disabled, respectively. We can then write the condition for optimal redistribution as

$$m^1 = m^2 \quad (1)$$

Note that this would not necessarily imply equal incomes. Because of differences in needs or tastes, the functional relationship between income and utility might differ between the workers and the disabled, but full equality is in any case a useful benchmark case.

Note that in this case *there is no trade-off between equality and efficiency*. If the market system were such as to lead to an efficient use of resources in the absence of taxes and transfers, this would still be the case after incomes had been redistributed. The point is that no prices are being distorted, so that the tax system does not create any efficiency loss. This does not mean that nothing has changed. The redistribution of income has shifted the competitive economy to another equilibrium, but this is efficient just as the initial equilibrium was.

Lump sum taxes are not very realistic as a description of real life policy instruments. One reason is, as emphasized by Mirrlees (1971), that the government does not have the information necessary to implement such taxes in an optimal fashion.⁷ We therefore move our analysis a step forward in terms of realism by assuming that taxation of workers takes the form of an income tax. In contrast to the lump sum tax the income tax has both income and substitution

⁷ Although Mirrlees (1971) is the classic reference in this area, the explicit reference to the informational problem came later, both in his work and others'. Several years earlier the problem of optimal income taxation had been formulated very clearly by Vickrey (1945) in an article that in several respects was much ahead of its time. For an informal discussion of the work of Mirrlees and Vickrey see Sandmo (1999).

effects; like the lump sum tax it reduces disposable income, but in addition it also distorts relative prices in favour of leisure and household work and against market labour supply. It is the latter effect which leads to a loss in terms of efficient use of resources, because a wedge has been driven between the worker's marginal productivity and his subjective evaluation of leisure. This efficiency loss is the marginal cost of redistribution, and because of this cost it will now be rational to carry redistribution less far than under the hypothetical lump sum tax. It turns out that the marginal efficiency cost can be expressed as the product of the marginal tax wedge and the reduction in labour supply, so that the condition for optimal redistribution becomes

$$m^1 = m^2 (1 - \theta e) \quad (2)$$

where θ is the tax wedge⁸ and e is the workers' compensated elasticity of labour supply with respect to the after-tax wage, which is positive. Thus, redistribution stops short of the point of equality of marginal utilities of income and is limited by the magnitude of the elasticity of labour supply. If the two marginal utility of income schedules are the same, the optimum will occur at a point where $m^1 < m^2$, so that workers have the higher income. From (2) we also see that the higher the elasticity, the less is the optimal degree of redistribution. If the compensated elasticity were to be zero, the income tax is equivalent to a lump sum tax (having only income effects), and the optimal degree of redistribution in the two cases is the same.⁹

The marginal cost of redistribution has a wider interpretation as the marginal cost of public funds and is of importance not only for the redistributive policies of the government, but also for its supply of public goods like communications infrastructure and public safety and private goods like child care and higher education.

Although I have here expressed the efficiency loss from tax finance and the costs of redistribution as stemming from the personal income tax and the labour supply response, the point is of course a much more general one. All the major types of taxes, with the exception of taxes on goods that are harmful to the environment, involve price distortions and efficiency losses. Just as the efficiency loss from taxes limits the degree of redistribution, it also limits the efficient

⁸ The tax wedge is the percentage difference between the consumer and producer wage. If t is the marginal tax rate, the two are related by the equation $\theta = t / (1 - t)$

⁹ This analysis has been presented in more detail in Sandmo (1991). Note that by solving equation (2) for the tax rate, one obtains an optimal tax formula which is a special case of the famous inverse elasticity rule, whereby the tax rate should be higher, the lower is the compensated price elasticity. Note also that for the formula to make sense, it must be the case that $(1 - \theta e) > 0$. This condition is satisfied if tax revenue is an increasing function of the tax rate, so that we are on the upward-sloping part of the "Laffer curve".

supply of public goods. However, the efficiency loss from taxation stems in large measure from its redistributive function. Although marginal tax rates involve a social cost, they also reflect a possible gain in the form of a more equitable distribution of income. Both aspects should be taken into account in calculations of the optimal size of the public sector.¹⁰

5. Concluding Remarks

In this brief survey I have argued first that it is by no means a trivial matter to decide how much redistribution the welfare state actually achieves, since this involves comparison with the unobservable situation where the state does not engage in redistribution. I have also pointed out that much of the public redistribution scheme consists in redistribution between periods and states in a person's life that might conceivably have been achieved through private markets; nevertheless, there are several arguments for having the state carry out some of this income smoothing on behalf of individuals. Finally, I have argued that the amount of redistribution between persons should, in a rational calculation of costs and benefits, be limited by the efficiency losses arising from the tax distortions.

There are some problems that I have not touched upon. Let me just mention a few. As Assar Lindbeck (1995) has recently pointed out, conventional economic analysis assumes that incentives can be changed without any effect on preferences, whereas it would be more realistic to assume that in the long run, the economic system will have an effect on the preferences of the members of society. Thus, a weakening of the economic incentives to work might over time strengthen people's preferences for not working, through a softening of the social sanctions against idleness. On the other hand, a society with a low degree of economic and social stratification might be less prone to social conflicts and therefore more productive. Clearly, the assumption that preferences are independent of the economic system is a strong one, and further research is needed to explore more general models.

I have also neglected a whole set of problems suggested by political economy considerations. The individuals in society do not simply adapt passively to the economic system. As political agents they also try to change the system to their own advantage. In a society where one's disposable income to a large extent is determined by government policy, much private effort will be directed towards changing the rules of the redistributive game. This effort has an alternative cost in terms of productive work and should clearly be counted as one of the

¹⁰ I have discussed this in more detail in Sandmo (1998).

additional costs of running a redistributive society.

While in the present paper I have mainly taken a normative approach to the problems of the welfare state, there is also a lot to be said for the more positive approach, whereby the goal is to explain the policies actually chosen, rather than point out the optimal ones. A positive study of the development of the welfare state would have to consider both the choice of institutional framework and the choice of policies, given the institutions. A fertile meeting ground for positive and normative analysis will be in the study of how institutions influence outcomes and the degree to which actual outcomes conform to the optimality conditions of normative economic theory.

References

- Akerlof, George. 1970. "The market for lemons: Quality uncertainty and the market mechanism," *Quarterly Journal of Economics* 84, 488–500.
- Andreoni, James. 1990. "Impure altruism and donations to public goods: A theory of warm-glow giving," *Economic Journal* 100, 464–477.
- Arrow, Kenneth J. 1953. "Le role des valeurs boursières pour la répartition meilleures des risques," *Econometrie*, CNRS, Paris. Translated as "The role of securities in the optimal allocation of risk-bearing," *Review of Economic Studies* 31, 1964, 91–96.
- Atkinson, Anthony B. 1991. "Social insurance," *The Geneva Papers on Risk and Insurance Theory* 16, 113–131.
- Barr, Nicholas. 1987. *The Economics of the Welfare State*, London: Weidenfeld and Nicolson.
- Debreu, Gérard. 1959. *Theory of Value*, New York: Wiley.
- Fritzell, Johan. 1991. *Icke av marknaden allena: Inkomstfördelningen i Sverige*, Stockholm: Swedish Institute for Social Research, Stockholm University.
- Harsanyi, John C. 1955. "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility," *Journal of Political Economy* 63, 309–321.
- Lindbeck, Assar. 1983. "Interpreting income distributions in a welfare state: The case of Sweden," *European Economic Review* 21, 227–256.
- Lindbeck, Assar. 1995. "Welfare state disincentives with endogenous habits and norms," *Scandinavian Journal of Economics* 97, 477–494.
- Lundholm, Michael. 1998. "Ex-post efficient social insurance," unpublished paper, Department of Economics, Stockholm University.
- Malinvaud, Edmond. 1972. *Lectures on Microeconomic Theory*, Amsterdam: North-Holland.
- Mirrlees, James A. 1971. "An exploration in the theory of optimum income taxation," *Review of Economic Studies* 38, 175–208.
- Okun, Arthur. 1975. *Equality and Efficiency: The Big Trade-off*, Washington, D.C.: The Brookings Institution.
- Persson, Mats. 1998. "Reforming social security in Sweden" in H. Siebert (ed.), *Redesigning Social Security*, Tübingen: Mohr Siebeck.
- Rawls, John. 1971. *A Theory of Justice*, Cambridge, Mass.: Harvard University Press.
- Sandmo, Agnar. 1991. "Economists and the welfare state," *European Economic Review* 35,

213–239.

Sandmo, Agnar. 1998. "Redistribution and the marginal cost of public funds," *Journal of Public Economics* 70, 365–382.

Sandmo, Agnar. 1999. "Asymmetric information and public economics: The Mirrlees-Vickrey Nobel Prize," *Journal of Economic Perspectives* 13, 165–180.

Sen, Amartya K. 1967. "Isolation, assurance and the social rate of discount," *Quarterly Journal of Economics* 81, 112–124.

Strotz, Robert H. 1955–56. "Myopia and inconsistency in dynamic utility maximization," *Review of Economic Studies* 23, 165–180.

Sugden, Robert. 1982. "On the economics of philanthropy," *Economic Journal* 92, 341–350.

Varian, Hal R. 1980. "Redistributive taxation as social insurance," *Journal of Public Economics* 14, 49–68.

Vickrey, William S. 1945. "Measuring marginal utility by reactions to risk," *Econometrica* 13, 319–333.