Distributional Objectives Should Affect Taxes but Not Program Choice or Design

Author(s): Aanund Hylland and Richard Zeckhauser

Source: *The Scandinavian Journal of Economics*, Vol. 81, No. 2, Measurement in Public Choice (1979), pp. 264-284
Published by: Wiley on behalf of The Scandinavian Journal of Economics
Stable URL: https://www.jstor.org/stable/3439965
Accessed: 27-05-2020 17:23 UTC

DISTRIBUTIONAL OBJECTIVES SHOULD AFFECT TAXES BUT NOT PROGRAM CHOICE OR DESIGN

Aanund Hylland and Richard Zeckhauser *

Harvard University, Cambridge, Mass., USA

Abstract

A society can redistribute income through the tax system, and through the choice and design of government programs. Neither type of redistribution is as efficient as lump-sum transfers would be, if feasible. In practice, however, both taxes and government programs serve redistributional goals. The question becomes how best to integrate them to achieve an optimal outcome, maximizing the redistributional effect for a given efficiency cost. The following conclusions are reached. If total benefits are independent of the income distribution and relative benefits are determined by before or after-tax income, those projects that yield the greatest total of unweighted benefits across the population should be selected. If benefits depend on the distribution of income, the optimal program will be one which produces maximal net benefits at the income distribution which is being induced. Redistribution is a concern, but is carried out solely through the tax system.

I. Introduction

A society can redistribute income through the tax system, and through the choice and design of government programs. Neither type of redistribution is as efficient as lump-sum transfers would be, if feasible. In practice, however, both taxes and government programs serve redistributional goals. The question becomes how best to integrate them to achieve an optimal outcome, maximizing the redistributional effect for a given efficiency cost.

The design of a society's tax schemes and other government programs should perhaps be thought of as a constrained maximization problem. It could take various forms—for instance, provide the minimum acceptable level of welfare to each citizen at the least cost to the citizens who would be providing resources. Or, given the political and financial constraint imposed by the willingness of non-poor citizens to redistribute funds, generate the maximum welfare level for the poor. On the other hand, a society may make these decisions as if it were maximizing some social welfare function and therefore making trade-

* Hylland's work is supported by a fellowship from the Norwegian Research Council for Science and the Humanities (NAVF); Zeckhauser's work is supported by NSF Grant SOC 77-16602 to Harvard University. Albert Nichols gave us helpful comments.

offs among the welfares of different groups. In any case, the pursuit of optimality in redistribution is critical to more fundamental normative investigations: Given some criterion of social welfare, derived perhaps from philosophical investigation, but in any case for our analysis assumed to be exogenously given, how should the potential institutions of society for transferring resources be arranged?

A fundamental issue in the literature on the expenditure side of public finance has been how to take account of distributional consequences in the choice and design of government programs. This problem is a subsidiary question in our more general analysis of optimality in redistribution. The central question we shall pursue is how to design taxes and government programs to maximize any arbitrary criterion for social welfare. This formulation is sufficiently general to include any of the problems mentioned above.

This paper adopts the approach of the optimal income tax literature in a number of respects. It assumes that individuals have differentiated ability levels which affect their opportunities for earning, that is, their wage rates. It assumes that each individual will react in a rational, self-interested manner to whatever system of taxes and government programs is enacted. Income, but neither leisure nor ability, can be observed and taxed. The formal model on which the optimal income tax theory is based is presented in Section II, where we also comment briefly on the theory and some results from the literature.

We have diverged from optimal income tax discussions in making government programs a major element of our model. To do so, we have assumed that benefits from such programs can be converted into an increased income equivalent. This increased income equivalent is computed assuming that an individual's money income is known.

The government programs we consider have distributional consequences; that is, benefits depend on income. This suggests as well that these programs have redistributional capabilities.

In our principal model we assume that the total of benefits all individuals receive from a particular program is independent of the income distribution in society. This total may depend on the ability distribution, which is exogenously given. For a particular program, the relative levels of benefits for various income groups are assumed to be known or derivable.

The critical question is how to design expenditure and tax programs in concert so as to maximize the social welfare criterion. Intuition might suggest that redistributional purposes should be pursued in both areas. That is, in addition to implementing a tax mechanism that promotes redistribution, a government program which is somewhat inefficient in the sense that it does not maximize unweighted net benefits might nevertheless be adopted because of its distributional effects. For example, given that we are pursuing distributional objectives, it might seem reasonable, when choosing among governmental programs, to attach different weights to the benefits going to different income groups.[1]

For our model, we prove this intuitive conclusion wrong. In the optimal arrangement, distributional objectives are achieved through the tax system alone. Government programs are chosen solely on the basis of efficiency criteria, that is, total net benefits are maximized. Individuals' benefits from the government programs are then taken account of, in a straightforward way, through design of the tax system. The proof of this result is not difficult and can be readily grasped, although the conclusion, we believe, runs contrary to the conventional wisdom in most liberal democracies.

It is often argued that we should redistribute through taxes rather than government programs because the latter entail great inefficiencies in the form of administrative costs or giving individuals goods they would not themselves have purchased. Whatever the merits of these arguments, they are irrelevant in our analysis, since we examine programs on the basis of net cash-equivalent benefits to individuals.

Our result holds regardless of whether the benefits an individual receives depend on income before or after taxes (Section V). There are various cases in which the arguments do not apply or only apply in part (Section VI); further study is needed here to see if related results can be obtained. These cases include those in which:

(*a*) Total benefits from a program depend on the income distribution.

(*b*) Benefits to an individual depend on ability, instead of or in addition to income.

(*c*) Benefits from programs are complementary to leisure, that is, the money-equivalent benefit may depend on the amount of leisure consumed as well as on income.

(*d*) Benefits depend on income in a non-deterministic way, so that there are differences in the benefits the members of an income group receive.

Our model does not consider a number of political factors which are important in real-life decision making with respect to redistribution. Some of these factors are discussed in Section VII.

II. The Theory of Optimal Taxation

The idea that the income tax scheme should be designed so as to maximize total social utility (or some more general function of individual utilities), is an old one. Musgrave (1959), Chapter 5, reviews the classical discussion,

[1] There is a rich literature within economics, indeed an entire subfield of the expenditure side of public finance theory, which pursues this mode of approach. Sometimes efficiency and distributional benefits are separated; other times total benefits are simply computed on a weighted basis. Much of this literature seems to be based on an assumption like the one mentioned in the text. (Other interpretations are possible, because the tax system need not be treated as a control variable.)

describing various criteria of social welfare, or equivalently, various ways of measuring the total sacrifice imposed by the tax levy. In general, this literature does not consider the possible effects of the tax system on people's choices of how much to work. The contribution of the modern optimal income tax theory has been to take this effect explicitly into account. Zeckhauser (1969) pursued this approach and solved the problem in a simple case. Mirrlees (1971) and Fair (1971) attacked the problem on a higher level of generality, and from their works an entire literature has emerged. We follow the tradition established in this literature in presenting the basic model.

Individuals are characterized by a single non-negative parameter, called ability and denoted a. The number a is the productivity of one unit of the person's labor. We assume perfectly competitive labor markets, so that the person is also paid a per unit of labor. Individuals know their own ability, but ability cannot be observed by the government and hence cannot be made the basis for taxation or other administrative decisions. The distribution of ability in the population is known; we assume that the distribution is absolutely continuous with density function f and finite expectation.[1] An individual derives utility from consumption goods and leisure. We assume that the effect on utility of all consumer goods can be captured by making after-tax income, denoted x, an argument of the utility function.[2] The effect of leisure is accounted for by making the amount of labor provided, y, an argument. Hence the utility function is of the form $u(x, y)$; it is the same for all individuals and is defined for all $x>0$ and $0 \leqslant y < 1$. (The restriction on y amounts to choosing the unit of labor such that 1 is the physical maximum. When the unit of ability is chosen, the unit of income is then given.) The function is assumed to be strictly increasing in x, strictly decreasing in y, strictly concave and continuously differentiable; $\lim_{x\to 0+} u(x, y) = -\infty$ for all y and $\lim_{y\to 1-} u(x, y) = -\infty$ for all x.

The government raises revenue by taxing individuals. An individual's gross income is the only variable which can be observed by the government and which therefore can be made the basis for taxation. This income can be observed without error. Let T be the tax scheme; then a person with ability a who provides y units of labor has a gross income of ay and pays a tax of $T(ay)$. We write $z = ay$ for gross income and allow $T(z)$ to be positive or negative. A negative value of $T(z)$ represents a welfare grant or subsidy.[3]

---

[1] The continuity assumption is made in order to simplify notation and is not essential.

[2] This assumption clearly holds in an economy with only one private good. For economies with many private goods, the utility function in our model must be interpreted as an indirect utility function; it is based on the assumption that people chose an optimal bundle of private goods for any level of after-tax income. An alternative model, not considered here, would have the utility function depend directly on consumption of the various goods.

[3] A negative $T(z)$ for low values of z is a feature of many existing tax and welfare systems. It is usually referred to as a welfare grant or the like, but could as well be called a "negative income tax", a term that seems more natural in the context of our model.

When the tax schedule is given, the utility-maximizing individual with ability a will face the following problem:

$$\text{Find } y \quad \text{with } 0 \leqslant y < 1 \quad \text{to maximize } u(ay - T(ay), y). \tag{1}$$

The way we have expressed the utility function, it is clear that all after-tax income will be consumed; hence y is the only decision variable. If the function T satisfies some weak conditions, this maximization problem will always have a solution.[1] We will also assume that the solution is unique. Let y_a be the optimal value; it obviously depends both on a and T. The optimal before-tax and after-tax income and utility level will be denoted $z_a = ay_a$, $x_a = z_a - T(z_a)$, and $u_a = u(x_a, y_a)$.

Revenue is required for programs outside the tax system. The revenue requirement is exogenously given and equal to R. Presumably, $R > 0$; it is not necessary, however, to assume this. (One can imagine the government having other sources of income, so that the income tax system can be allowed to run a deficit.) T must be chosen so that the net revenue from the income tax system is at least R, that is

$$\int_0^\infty T(z_a)\, f(a)\, da \geqslant R, \tag{2}$$

where z_a depends on a as described above.

The government's objective is represented by some criterion of social welfare, which depends on everybody's utility level. For simplicity of notation, we assume that the criterion can be expressed by some social welfare function; hence the government's objective is to

$$\text{maximize } \Phi(u_a;\, a \geqslant 0) \tag{3}$$

for a given function Φ. Note that the argument of Φ is the infinite-dimensional vector of numbers u_a for $a \geqslant 0$. Later, we will write $\bar{u}$ for this argument.[2] The formulation (3) is quite general, but it does imply that social welfare depends only on individuals' utilities and abilities. It does not depend directly on income and amount of labor provided. We expect Φ to be monotone, that is, $\Phi(\bar{u}) \geqslant \Phi(\bar{u}')$ if $u_a \geqslant u'_a$ for all $a \geqslant 0$. (Not all arguments below depend on monotonicity of Φ, however.)

In the discussion below, we are not actually going to compute optimal tax schemes; hence we may as well keep the general formulation (3) of the social welfare function. Some comments about possible forms of the functions are, however, in order.

---

[1] Continuity of T will suffice; for a weaker sufficient condition, see Mirrlees (1971) p. 177.
[2] Formally, the argument is a function from the set of non-negative real numbers into the real numbers; thus Φ itself is a functional. Of course, the criterion of social welfare should be allowed to depend not only on the numbers u_a, but also on how many people have each ability level a. Since the latter is exogenously given (by the function f), it can be incorporated into the functional form Φ.

For one thing, the different versions of constrained maximization problems mentioned in the introduction can be expressed in this model. One possibility is to construct Φ so that it has negative values when the constraint is not satisfied and non-negative values when it is satisfied, while it otherwise represents the chosen social welfare criterion. (In general, this implies that Φ will have a discontinuity corresponding to the constraint.) Alternatively, we can carry an explicit constraint through the entire argument below. It should be noted that the introduction of a constraint, in either of the two formulations, adds to the computational problems of actually finding optimal tax schemes, but does not raise any basic conceptual issues.

Possible Incorporation of Altruistic Concerns. Our assumptions concerning the form of the utility function u imply that everybody is completely selfish. This assumption, however, is not essential for the subsequent discussion. Under an alternative interpretation, u can be viewed not as a utility function capturing everything which is relevant to the individual, but merely as an index of personal satisfaction. The person's utility is then given by $v = v(u_e, \bar{u}_{-e})$, where u_e is the person's own level of satisfaction, and $\bar{u}_{-e}$ is the vector of these levels for everybody else. The function v specifies the degree of selfishness; one extreme case is given by $v(u_e, \bar{u}_{-e}) = u_e$, at the other extreme, u_e contributes to the functional value in exactly the same way as any component of $\bar{u}_{-e}$. It is assumed that the functional forms of u and v are the same for everybody, and that v is increasing in u_e and treats the components of $\bar{u}_{-e}$ symmetrically. This does not mean that the level of satisfaction of everybody else in any sense must be given equal weight; it is quite possible, for example, to pay more attention to the less well off. But the identity of other individuals cannot be taken into account. Neither can v depend directly on the consumption of goods and leisure of other people; only their satisfaction levels matter.

In this setting, the rational individual will still choose y according to (1). This is so because the individual in no way controls $\bar{u}_{-e}$; hence maximizing v is equivalent to maximizing u. The social objective would be to

$$\text{maximize } \Psi(v_a; a \geqslant 0) = \Psi(\bar{v}) \tag{3'}$$

for some function Ψ, where v_a is the utility level achieved by a person of ability a when everybody acts according to (1). The number v_a depends on the entire vector $\bar{u}$. The vector of numbers v_a for $a \geqslant 0$ is denoted $\bar{v}$.

The formulation (3′), however, is no more general than (3). For given v and Ψ, one can simply set $\Phi(\bar{u}) = \Psi(\bar{v})$, which is well-defined by the symmetry assumption we have imposed on the function v. Monotonicity of Φ implies some restrictions on the functions v and Ψ; these will, for example, be satisfied if Ψ is monotone in $\bar{v}$ and $v = v(u_e, \bar{u}_{-e})$ is monotone in $\bar{u}_{-e}$. The latter condition rules out such possibilities as $v(u_e, \bar{u}_{-e})$ depending on the relative position of u_e among the components of $\bar{u}_{-e}$.

The Expected Utility Approach or Utilitarianism. Classical utilitarianism corresponds to

$$\Phi(\bar{u}) = \int_0^\infty u_a f(a)\, da \tag{4}$$

This is the criterion by which people would evaluate tax rules if they were ignorant of their own ability and knew only the probability distribution f, provided that u really captures everything which matters to the individual and represents attitude towards risk. That is, u must be a von Neumann–Morgenstern utility function. (Elsewhere in the paper we need only assume that u is a value function representing preferences under certainty.) This contractual formulation does not correspond to any real-world decision-making situation. In practice, at the time decisions are made, much of the uncertainty about an individual's ability has been resolved. Assuming self interest, those who thus far have been fortunate will favor a less progressive tax scheme and vice versa. A question of importance for both policy and philosophical investigation is: To what extent should arguments about this hypothetical "state of ignorance" influence real-world decision-making?[1]

Preference for Equality. The utilitarian formulation does not rule out a preference for equality; such a preference emerges if u exhibits risk aversion in income and hours of work. What is ruled out, is a desire for equality over and above what is implied by risk aversion. Such an additional preference for equality (or for more general distributional criteria) is captured by using the formulation involving the function v. The social welfare function is then given by

$$\Psi(\bar{v}) = \int_0^\infty v_a f(a)\, da. \tag{4'}$$

This is a special case of (3′) and therefore of (3), but it is more general than (4). In particular, (5)–(7) below can be obtained from (4′) by appropriate choice of v.[2]

If we want to promote equality as such, whether justified by an argument like the one behind formula (4′) or in some other way, we should pay more attention to the u-values of the less fortunate. This can be achieved by using a social welfare function of the form

$$\Phi(\bar{u}) = \int_0^\infty u_a g(a) f(a)\, da, \tag{5}$$

where g is a positive and decreasing weighting function,[3] or

$$\Phi(\bar{u}) = \int_0^\infty h(u_a) f(a)\, da, \tag{6}$$

[1] See Zeckhauser (1974) for some discussion of this issue.
[2] Harsanyi (1977), Chapter 4, argues that social welfare functions should always be of the form (4′).
[3] It is easy to see that $a > a'$ implies $u_a > u_{a'}$; therefore, giving more weight to individuals with low ability implies giving more weight to the utility of the less fortunate.

where h is increasing and concave. A limiting case of (5) and (6) is the maximin rule, given by

$$\Phi(\bar{u}) = u_0, \tag{7}$$

provided that there exist individuals with ability arbitrarily close to 0.[1]

The Solution. When the social welfare function has been specified, the government's problem is: For given u and R, find the tax scheme T which maximizes (3) subject to (1) and (2). This is not at all a trivial problem, as the literature on the subject clearly shows. In general, we do not even know that an optimal solution exists.[2] In a sense, our discussion below presupposes that this problem has been solved. But this is not as important a restriction as it may seem. Provided that there is an upper bound on the achievable values of Φ, our results will essentially hold even if one can only find tax schemes which approximate the upper bound on Φ, which is likely to be an easier problem. See discussion at the end of Section III below.

The model just described makes strong simplifying assumptions. In addition to more formal simplifications, some of these are: Differences in tastes are ignored. The population is fixed; hence in and out migration is assumed to be impossible. Income can be perfectly observed, and the cost of administering the system is independent of the tax scheme. The time frame is ignored, and no attention is paid to the problem of defining the consumption unit. Finally, and perhaps most importantly, the work/leisure choice is assumed to be a pure problem of utility maximization, and productivity and wage rates are independent of the choices people actually make. This rules out, for example, any kind of institutional constraints such as standard working hours, or a feedback between hours worked and productivity.

Therefore, conclusions drawn from the model should not be interpreted as firm policy recommendations, but rather as indications of what an optimal solution might look like and how it will depend on the parameters. For further discussion of the problems and for a number of results in special cases, we refer to Mirrlees (1971), Fair (1971), Atkinson (1973) and Feldstein (1973).

III. Government Programs with Income-related Benefits

The general question we want to ask is: What is the optimal simultaneous choice of tax schemes and government programs when the latter have distributional effects? In this section, we address the simpler problem of designing the

[1] In general, we can define $a_0 = \inf \{a \mid \int_0^a f(\alpha)\,d\alpha > 0\}$. The number a_0 is then essentially the lowest existing ability, and the maximin social welfare function is $\Phi(\bar{u}) = u_{a_0}$.

[2] Mirrlees considers social welfare criteria of the form (6), of which (4) is a special case. From relatively weak conditions on the utility function u, he succeeds in proving the existence of an optimal tax scheme and deriving some general properties. But more specific results are obtained only when particular forms of u are assumed. Other authors have simplified the problem by restricting T to particular functional forms. Thus they derive, for example, the optimal linear tax scheme.

tax scheme when the government program is given. Here we assume that the benefits people derive from the program depend on before-tax income. The case of benefits depending on after-tax income is technically a little more complicated, and is considered in Section V.

A government program P is characterized by: B, a real number, representing the total (monetary) benefits from the program; C, a real number, representing the costs to the government of implementing the program; β, a function, defined for all $z \geq 0$, such that $\beta(z)$ represents the relative benefit from the program to a person whose before-tax income is z, as described in eqs. (8) and (9) below. (Since β represents relative benefits, nothing is changed if all values of β are multiplied by some positive constant.)

The entities are supposed to represent net benefits and costs; hence we can include in the model activities which are partially financed by user fees or the like. Our terminology might be thought to imply that B, C and $\beta(z)$ are all non-negative, but nothing in the formal derivations requires that this be the case. Therefore, we can also include programs which save money for the government by imposing income-related costs on individuals.[1]

We have defined β as a *relative* benefit function since we assume that B, the total program benefits, is fixed. If $\beta(z)$ were defined as the absolute value of benefits to a person with gross income z, the implication would be that the total benefits produced by the program would depend on the income distribution; in particular, if β has a maximum at z_0, total benefits could be increased by more people earning z_0. (Note that the income distribution is endogenous, it is determined by individual optimization according to (1) when the tax scheme is given.) We assume instead that the total benefits to all individuals are constant and independent of the income distribution. That is, we assume that some kind of a divide-the-spoils or congestion effect occurs if there is an increase in the number of people in the income group which receives the highest relative benefits from the program. (A more general formulation is discussed in Section VI.)

In an individual's utility, the benefits from the program are supposed to have the same effect as an increase in income proportional to $\beta(z)$, where z is before-tax income. (It does not matter whether the benefit is added to before-tax or after-tax income.) If z_a is the before-tax income of a person with ability a, that person's benefits will be

$$b(z_a) = b_0 \beta(z_a), \tag{8}$$

where b_0 is a number which satisfies

$$\int_0^\infty b_0 \beta(z_a)\, f(a)\, da = B. \tag{9}$$

[1] The military draft can perhaps be seen as such a program.

When a tax scheme T' and the function b of benefits are given, an individual with ability a will choose the amount of work to provide by solving the problem[1]

$$\text{Find } y \quad \text{with } 0 \leqslant y < 1 \quad \text{to maximize } u(ay - T'(ay) + b(ay), y). \tag{10}$$

As before, let y_a be the solution to this maximization problem for a given a, and let z_a, x_a and u_a be optimal before-tax income, after-tax income and utility. If the government's revenue requirements for other purposes than the program under consideration are R_0, the tax scheme must be chosen so that

$$\int_0^\infty T'(z_a) f(a)\, da \geqslant R_0 + C. \tag{11}$$

When the tax scheme T' is given, the benefit function and the income distribution will be mutually dependent on each other through eqs. (10) and (8)–(9); hence they must be determined simultaneously. When they are computed, we can check whether (11) holds. Alternatively, we can say that the functions T' and b, the constant b_0 and the income distribution must be chosen simultaneously so as to satisfy (8)–(11).

As before, the objective is to maximize a certain social welfare function given by (3) or one of the special forms (4)–(7). The achievable levels of social welfare with and without the program P, are related in the following way.

Proposition 1. Connection between tax schemes in the presence and absence of programs

Let u, B, C, β and R_0 be as described above. Suppose that T is a tax scheme such that in the absence of P, (2) is satisfied with $R = R_0 - B + C$. Then there exists a scheme T' which satisfies (8)–(11) such that every individual reaches equal utility levels when T' is used and P is implemented and when T is used and P is not implemented. Conversely, if T' satisfies (8)–(11), there exists a T which raises at least $R_0 - B + C$ in revenue when P does not exist, such that T without P and T' with P leads to the same utility level for everybody.

Proof

Let T be given. Via (1), T induces an income distribution $\bar{z} = (z_a;\, a \geqslant 0)$. Compute b_0 and b from this distribution by (9) and (8). Define, for all $z \geqslant 0$,

$$T'(z) = T(z) + b(z). \tag{12}$$

Now (1) and (10) are exactly the same expressions for all values of a and y. Hence the solution is the same for each a, and T and T' will induce the same

[1] The formulation assumes that the individual regards the function b as constant, although in fact b depends on people's behavior via (8) and (9). This corresponds to the usual assumption in economics of price-taking or competitive behavior, an assumption which is reasonable when any one individual's action has only a negligible impact on society.

distribution of gross income. Therefore, (8) and (9) hold when the income distribution resulting from T' is used. By assumption, (2) holds with $R = R_0 - B + C$; hence (9) and (12) imply (11). Finally, it is clear from (1) and (10) that the resulting utility level u_a is the same under the two regimes, for any a.

Conversely, assume that T' satisfies (8)–(11) for some constant b_0 and some function b. Define $T(z) = T'(z) - b(z)$. Again, expressions (1) and (10) become equal, and an argument similar to the one used above will apply. The proof is complete.

Hence the problem of finding an optimal tax scheme in the presence of a government project with income-related benefits is reduced to the corresponding problem in the absence of such programs. As has been pointed out earlier, the latter problem is non-trivial. If the set of achievable values of Φ is bounded from above, it may be substantially easier to find tax schemes which approximate the least upper bound of the set. By Proposition 1, the set of achievable values of Φ must be the same in the situation with the program and in the appropriate situation without it, and an approximation of the upper bound in the latter case can immediately be transformed into an equally good approximation in the former case.

Then one can ask whether such approximations represent a satisfactory solution to the optimization problem, and whether the condition that the achievable range of Φ be bounded from above is an important restriction. (This question applies equally well to the original optimal income tax problem as to our extension of it.) If we view Φ solely as a representation of our ordinal preferences on social utility distributions, then this restriction is vacuous; it is always possible to find an order-preserving transformation of Φ which makes it bounded. But if Φ is interpreted this way, it is not at all clear that approximating the upper bound on Φ in any real sense implies coming close to an optimal solution. We would, on the other hand, like to interpret Φ as some kind of cardinal measure of social welfare, however vaguely that concept might be defined. Then the boundedness condition follows from assuming that society's resources and ability to achieve its goals are limited, an assumption probably accepted by most people. Under this interpretation, one can reasonably claim that approximate solutions of the type considered here are satisfactory.

IV. Comparison of Alternative Programs

Now assume that P_1 and P_2 are two programs of the type described in the previous section, characterized by numbers and functions B_1, C_1, β_1 and B_2, C_2, β_2, respectively. The functions β_1 and β_2 can be different, hence the two programs can distribute the benefits in widely different ways among income groups. The general form of our social welfare function (3) allows us to differentially evaluate benefits to different income groups; in particular, we are allowed to put a higher weight on benefits to people who are relatively worse off. Our major result is that in spite of this, if optimal taxation is available we should choose between competing government programs P_1 and P_2 solely on basis of their net benefits B_1-C_1 and B_2-C_2.

Proposition 2. Comparison of alternative programs

Let programs P_1 and P_2 be given, and assume $B_1-C_1 \geqslant B_2-C_2$. Then the optimal social welfare level which can be achieved under P_1 is at least as high as the optimal level under P_2.

Proof

This is immediate from Proposition 1. Suppose that a certain level ϕ of social welfare can be achieved under P_2, and let T'' be the corresponding tax scheme. By the second half of Proposition 1, a tax scheme T exists which satisfies (2) with $R=R_0-B_2+C_2$ and which, in the absence of both programs P_1 and P_2, gives all individuals the same utility level as they get in the situation with P_2 and T''. Hence the social welfare level when T is used is ϕ. By assumption, $R_0-B_1+C_1 \leqslant R_0-B_2+C_2$; hence T also satisfies (2) with $R=R_0-B_1+C_1$. The first half of Proposition 1 then implies the existence of a tax scheme T' which, when P_1 is implemented, produces social welfare level ϕ. The optimal level under P_1 is therefore at least ϕ. The proof is complete.

In fact, we have proved something which is stronger than the statement of Proposition 2, namely the following: Let $B_1-C_1 \geqslant B_2-C_2$, and let $\bar{u}''$ be any vector of individual utility levels which can be achieved under P_2 and some tax system T''. Then there exists a tax system T' such that the vector of individual utility levels becomes $\bar{u}'$ under P_1 and T', and $u'_a \geqslant u''_a$ for all a.

Formally, Proposition 2 only considers the comparison between two alternative programs. But the result implies the existence of a consistent way of ranking mutually exclusive programs. An optimal decision rule will simply be: Choose the program with highest net benefits. Alternative ways of designing what is basically the same project can formally be seen as different programs; therefore, the result also implies that when designing a project, the configuration which maximizes net benefits should be chosen. The alternative "no program" can be viewed as one element of the set of alternative programs; it is characterized by $B=C=0$ and β arbitrary. If we are presented with a set of potential government programs which are not mutually exclusive, we can let every technically feasible subset of this set be a "program" in the sense of our model. This allows for complementarities in benefits and costs among the original programs, as long as all composite programs fit the model of Section III.

One can ask whether a program with higher net benefits actually leads to a higher achievable social welfare level. This is equivalent to asking whether a relaxation of the constraint (2) in the ordinary optimal income tax problem leads to a strict increase in the optimal value of (3). Under reasonable conditions on u and Φ this will be the case. We will not state and prove any formal result to this effect but only argue informally that it is likely to be true: Let T satisfy (2) for a given R, and let $r>0$ be the amount by which the revenue requirement is reduced. For some small number $t>0$, define T^* by $T^*(z)=T(z)-t$ for all z. If T^* is substituted for T and people do not change the amount of labor they provide, the revenue loss is t. Work decisions will in fact change, and this can increase the revenue loss. If individual decisions are continuous in t, the revenue loss is also continuous, and the loss can be restricted to r by choosing t small enough but positive. We assume that it is possible to choose t so that it depends only on r and not on T, at least as long as T is optimal or almost optimal. (This amounts to a regularity condition on u.) Everybody's utility has increased because of the change from T to T^*, and the increase is at least equivalent to a lump-sum monetary transfer of t. For all reasonable social welfare criteria, this leads to an increase in the value of Φ; in fact, it leads to an increase of at least ε, where $\varepsilon>0$ depends only on t. (The latter, stronger statement holds for all the special forms (4)–(7).) If now T is chosen so that the value of Φ is closer than ε to the upper bound when the revenue requirement is R, then T^* demonstrates that the social welfare level increases strictly when the requirement is reduced to $R-r$.[1] Note that we have not assumed that there actually exist optimal tax schemes; we have only assumed that the optimum can be approximated.

Even if the argument of the previous paragraph fails, our main result is still true. Proposition 2 implies that maximizing net benefit is *an* optimal decision rule, though it need not be the only optimal rule.

V. Programs Whose Benefits Depend on After-Tax Income

In the model presented above, relative benefits from a program depend on before-tax income; alternatively one can assume that they are determined by after-tax income. Formally, this amounts to a change in eqs. (8)–(10) in Section III. If s_a is the after-tax income of a person with ability a, when a certain tax scheme T' is used and P is implemented, eqs. (8) and (9) are replaced by

$$b(s_a)=b_0\beta(s_a), \tag{13}$$

where b_0 satisfies

$$\int_0^\infty b_0\beta(s_a)\,f(a)\,da=B. \tag{14}$$

[1] This does not imply that the across-the-board tax cut represented by T^* is the optimal response to a reduction in required revenue; it is just one possible response which will increase social welfare.

A person with ability a must solve the following problem, which corresponds to (10):

$$\left.\begin{array}{l}\text{Find } y \quad \text{with } 0 \leqslant y < 1 \quad \text{to maximize} \\ u(ay - T'(ay) + b(ay - T'(ay)), y).\end{array}\right\} \tag{15}$$

In the same way as before, (13)–(15) and (11) should be viewed as a set of conditions which must be simultaneously satisfied by T', b, b_0 and the income distribution s_a.

Provided that β satisfies certain regularity conditions to be specified below, a result similar to Proposition 1 can now be proved. An equivalent of Proposition 2 then follows directly.

To outline the proof, let T satisfy the premise of the first half of Proposition 1. T induces a certain distribution of after-tax income, given by x_a for $a \geqslant 0$. For a given number $b_0 \geqslant 0$, find s_a for $a \geqslant 0$ such that

$$s_a + b_0 \beta(s_a) = x_a. \tag{16}$$

If β is continuous and bounded, this equation always has a solution. The number s_a will represent after-tax income in the presence of P; therefore, we would like s_a to be non-negative for all x_a and b_0 that actually occur. (Formally, we can permit s_a to be negative, provided that β is defined on negative arguments.) Under any optimal or almost optimal tax scheme, there will be a positive lower bound on after-tax income. That is, for any T we want to consider, $T(0)$ is a non-negligible negative number, and $x_a \geqslant -T(0)$ for all a. Hence the condition $s_a \geqslant 0$ is not very restrictive. Also, we would like the solution of (16) to be unique. (On the formal level, this is not essential either; if there are several solutions, we just choose one of them.) This amounts to requiring that β not decrease too fast; in particular, if β is differentiable we must have $\beta'(s) > -1/b_0$ for all s and all b_0 which are being considered. (Intuitively, this is equivalent to saying that the benefits from the program should not fall so fast as income increases that the benefits lost outweigh the income gained.)

For any b_0, we now compute

$$\int_0^\infty (x_a - s_a) f(a)\, da.$$

For $b_0 = 0$, this is equal to 0. Under the conditions outlined above, the expression is an increasing function of b_0. Now we find b_0 such that

$$\int_0^\infty (x_a - s_a) f(a)\, da = B. \tag{17}$$

Such a b_0 will exist, provided that β is positive over a non-negligible range and B is not too large compared to the aggregate after-tax income. By the above, (17) determines b_0 uniquely.[1] Note that the larger we have to choose b_0, the more restrictive are the conditions discussed above.

When b_0 is determined and s_a is defined by (16), we define T' such that, for all a

$$T'(z_a) = z_a - s_a. \tag{18}$$

Here z_a is the before-tax income of a person with ability a who acts according to (1). If $z \neq z_a$ for all a, we let $T'(z)$ be some large number, for example, $T'(z) = 2z$. (The point is that nobody shall want to have gross income z.)

Now it is easy to see that (1) and (15) have the same solution y_a for every a. The other properties of T' required in the conclusion of Proposition 1 are also established in a straightforward manner.

The proof of the second half of the Proposition is less complicated and does not require extra assumptions. Let T' satisfy the premise. Then a function b and a number b_0 are also given, such that (13)–(15) and (11) hold. We define

$$T(z) = T'(z) - b(z - T'(z)), \tag{19}$$

and the conclusion follows immediately.

VI. Problems for Further Study

In this section, we present a number of cases in which the arguments of Propositions 1 and 2 do not apply or apply only in part. In these situations, it is possible that redistributional objectives should affect the choice and design of programs and not only the construction of the tax system. Whether and to what extent this will be true should be the object of further study.

Total Benefits Depend on the Income Distribution. We have assumed that the total benefits derived from a program are independent of the endogenously determined income distribution. More generally, one could have total benefits depend on this distribution. This is equivalent to saying that benefits to an individual with income z are $b(z, \bar{z})$, where b is an arbitrary function and $\bar{z} = (z_a; a \geqslant 0)$ is the income distribution. When such a function b is given, one can compute the total benefits $B(\bar{z})$ given any distribution $\bar{z}$. We have considered the special case in which b is such that B becomes a constant function. Another special case has $b(z, \bar{z})$ depend only on z; this implies the absence of congestion effects or the like.

In this model, the problem of designing an optimal tax system, given certain government programs, can be reduced to the similar problem in the absence of such programs. That is, we have a result which in a sense is similar to Proposition 1. To be precise, we have the following: Let T' be a tax scheme

[1] The discussion so far has assumed $B > 0$ and $\beta(s) \geqslant 0$ for all s. The case $B < 0$ and $\beta(s) \leqslant 0$ for all s can also be taken care of. But we cannot allow β to change sign.

which is feasible in the presence of a program P and leads to income distribution $\bar{z}$. Define $T(z)=T'(z)-b(z,\bar{z})$. If T is used and P does not exist, the income distribution will again be $\bar{z}$. Moreover, the revenue generated will be at least $R_0-B(\bar{z})+C$, and everybody's utility level will be the same as in the presence of T' and P. Conversely, if T leads to income distribution $\bar{z}$ and revenue no less than $R_0-B(\bar{z})+C$, we can construct a scheme T' which is feasible when P exists and such that the same kind of utility equivalence holds. Hence finding an optimal income tax when P is implemented is equivalent to solving the problem of Section III with the revenue constraint (2) replaced by

$$\int_0^\infty T(z_a)\, f(a)\, da + B(\bar{z}) \geqslant R_0 + C.$$

But it does not follow that the solution T in itself is an optimal tax scheme for any particular level of the revenue requirement R.

Programs cannot be compared directly on the basis of the net benefits $B-C$, as was done in Proposition 2, since B is not fixed. But a similar result does obtain. Assume that a program P_2 is implemented together with a tax scheme T'', such that the revenue constraint is satisfied and the income distribution $\bar{z}$ is induced. Then assume that there exists a program P_1 such that $B_1(\bar{z})-C_1 \geqslant B_2(\bar{z})-C_2$. That is, assume that P_1 has at least as high net benefits as P_2, when measured at the income distribution induced by P_2 and T''. Then we can prove, by an argument similar to the proof of Proposition 2, that there exists a tax scheme T' such that the combination P_1 and T' is at least as good, according to Φ, as P_2 and T''. If $B_1(\bar{z})-C_1 > B_2(\bar{z})-C_2$, it will normally be strictly better; see the discussion at the end of Section IV. The optimal configuration of a tax scheme T and a program P must therefore have the property that P is the program which maximizes net benefits, when benefits are measured at the income distribution resulting from P and T.

Benefits Depend on Ability. Next we consider the case in which total benefits from a program are constant, but individual's relative benefits depend not on income but on the unobservable variable ability (or that they depend both on income and ability). The proof of Proposition 1 cannot be applied. The tax system T' was constructed so as to "tax away" all benefits from the program, thereby eliminating any distributional effect. This is impossible when benefits depend on ability.

Normally, we would expect income to be a strictly increasing function of ability, provided that the tax system is optimal and people act rationally.[1]

[1] In the lower part of the ability range this cannot be expected to hold; under reasonable social welfare functions and optimal taxation there will exist a constant $a_0>0$ such that individuals with ability less than or equal to a_0 do not work and hence have the same before-tax income. (This is proved by Mirrlees in his model.) For ability level above a_0, if income is not strictly increasing in ability, income plays a role similar to that of a Giffen good. This possibility does not contradict our assumptions, but is certainly something out of the ordinary.

Then ability can be inferred from income, and one can ask whether that fact can be used to obtain a result equivalent to Proposition 1. The answer is no, for the following reason: Let a tax scheme T be given, as in the proof of the Proposition. Since ability can be inferred from income, one can construct a function b such that $b(z)$ is the benefit received from the program P by a person with income z, provided that the income distribution is the one induced by T. Then T' can be defined by (12). The expression the individual will maximize is not (10) but $u(ay - T'(ay) + b'_a, y)$, where b'_a is the benefits which accrue to a person with ability a. This is not equivalent to (1), and the proof breaks down. By working a little more or a little less, an individual will be perceived by the tax scheme as having a little higher or lower ability. This influences the "benefit part" $b(z)$ of the tax given by (12), but benefits are related to ability and do not change; hence the incentives are distorted.

Benefits Are Complementary to Leisure. The assumption that benefits are equivalent to an income-related increase in income, essentially rules out programs which are complementary to leisure. A way of removing this restriction is to let benefits depend on the amount of leisure consumed. But leisure is an unobservable variable, and we run into the same difficulties as we did above in considering ability-related benefits. For any level of gross income, leisure can be seen as a function of ability and vice versa. Hence the two cases are equivalent.

Benefits Differ among Members of an Income Group. Thus far, we have assumed that two individuals who are equal in income and other factors relevant to the model receive the same benefits from a program. More generally, and clearly more realistically, one could have benefits depend on income in a non-deterministic way. This can be incorporated into the formal model by assuming that for each $z \geq 0$, there is a known probability distribution of relative benefit levels received by individuals with income z. In an important special case there are, for each income group, only two possible benefit levels, namely 0 and a positive level. This corresponds to programs which do not reach the entire target group, but which benefit equally all individuals with the same income who actually participate. Rate of participation and benefits to participants can depend on income.

In this model, a program has two kinds of distributional effects, corresponding to differences in benefits within and among income groups. The income tax scheme can in no way be used to compensate for differences of the first type; therefore, our previous arguments do not apply. Suppose that the social welfare function implies a preference for equality. (This will be the case, for example, if we use the expected utility formulation (4) and the utility function u displays risk aversion.) Other things equal, we would then prefer a program for which the differences in benefits within income groups are small. That is, we would be willing to make a sacrifice in total benefits in order to achieve greater homogeneity within income groups. In particular, we would accept lower average benefits for any one group if they were more evenly spread among the members of that group. More definite statements about this type of trade-offs can only be made if further assumptions are introduced concerning the programs, the utility function and the social welfare function. This is an area for further study.

Next we turn to the differences among income groups. One would perhaps expect that these differences could be eliminated through the tax system and therefore should not influence the choice of program, exactly as in our main model. In fact, the situation is more complicated. When there are differences in benefits within a group, risk is imposed on the members of that group. (Equivalently, they are subjected to variation in utility level.) Everybody is not equally able to bear risk, and this is a factor which should be taken into account when programs are designed and chosen.

Let us present an example: Suppose that the social welfare function is given by (4) while the utility function u displays decreasing absolute risk aversion in money. That is, everybody is risk-averse, but the rich are less so than the poor. Moreover, assume that every feasible program reaches half the population in every income group; the other half receives nothing. There is a choice between a program which concentrates the benefits in the upper end of the income scale and one which mainly benefits the poor; total net benefits are approximately equal for the two programs. Then we shall choose the former program, the one which mainly benefits the rich. This way we avoid placing any significant risk on the lower-income individuals, who are most strongly risk-averse. Instead, the risk is borne by people with higher income, who are better able to do so. The purely distributional aspect of the programs, that is, the fact that one of them directs the benefits towards the rich and the other one towards the poor, should not influence the choice between them. This effect is compensated for through the tax system, as in our main model.[1] Again, further study should be devoted to a detailed examination of more general cases.

In addition to the possibility that benefits from a program vary randomly within an income group, one can clearly imagine programs for which benefits depend on identifiable criteria other than income. Examples are programs which help victims of accidents or others who are "needy" in a sense not solely related to money. If the target group is well-defined and easily recognized, such a program obviously can achieve its goal in a more efficient manner than programs of the type considered earlier. In order to incorporate such programs into the model, we must let the common utility function depend on

---

[1] To be precise, this is merely a sketch of an example. We have assumed that people make their work decisions before they know whether they will benefit from the program. The risk they have to bear will affect these decisions, and a full analysis should consider this complication. The conclusion is unlikely to be affected, however.

other arguments than income and hours of work. No difficulties arise in our tax schemes if all of these other arguments are readily monitored and can be made bases for taxes.

VII. Political Aspects of Distributional Decisions

Our entire discussion has assumed that society is making one grand decision in which taxes and government programs are simultaneously determined. Hence the conclusions apply to a situation in which a constitutional contract is being designed, and to an ideal form of government which makes comprehensive decisions about all sides of government policy and is aware of and takes account of all interrelationships between different areas and activities. In these cases, the conclusion is clear: Distributional considerations should be taken into account when the tax system is designed and only then; therefore, political groups which have distributional objectives should focus their attention on the tax system. Conversely, if one is not satisfied with the level of redistribution which can be achieved through the tax system, programs of the type considered here cannot improve matters; one must look elsewhere.

Real-life politics is of course not like this. Government decisions are made one by one; they may influence each other, but not in the comprehensive way described above. If decisions about programs and taxes were completely independent, a group with distributional goals should pursue them in both areas. If there is an incomplete relationship between the two areas, our results suggest that the group should emphasize tax strategies, but other programs should not necessarily be neglected.

There are other features of the political system that may tend to diminish the relevance of our conclusions. We will not attempt to discuss this issue in any detail, but a few points will be raised.

For one thing, groups with distributional objectives will often find—or at least believe—that their goals can more easily be reached in one area than in another. For example, a group which works for increased well being for the poor may achieve greater success by urging subsidies for low-income housing than by advocating cash grants to the same low-income groups. That is, the former type of support may be more acceptable to the higher-income people who will have to pay the subsidy. This claim can be seen as an argument against the use of a social welfare function of the form (3). Social welfare, it can be argued, does not depend only on individual utility, but directly on the levels of individual consumption of various goods, at least as far as certain basic necessities are concerned. (A more general formulation which takes account of this possibility can still employ social welfare functions of the form (3′) or (4′), but the arguments of the utility function v must include everybody's consumption of the basic goods, or at least some measure of how these goods are distributed.) Indeed, it is often asserted that in some modern industrial societies citizens are more "goods egalitarians" than "income egalitarians". If this is true, elimination of direct transfer programs would have relatively little impact on the progressivity of the tax system and on balance would harm the poor; if it is wrong, significantly increased progressivity in the tax system would be accepted if transfer programs were abolished.

Moreover, some programs benefit identifiable groups, not defined by income but by some criterion which may be related to income, such as blindness or residence in a particular area. These people are not likely to be swayed by an argument that their income group as a whole would be better served by different programs or general transfers through the tax system. From the point of view of various subgroups of low-income people, economic transfers have the character of a public good. Our results indicate that the income class as a whole should prefer that the most efficient programs be adopted and transfers made through the tax system. But each subgroup will prefer that the particular program which benefits that group be implemented. The program may reduce the willingness of higher-income groups to make other transfers, but this effect is spread out over all low-income people, and the subgroup has made a net gain.

Finally, it should not be forgotten that those who provide government services have a say in the political process. This is yet another reason why the outcome is not always what our model of rational and simultaneous decision making predicts.

VIII. Concluding Remarks

The implications of this analysis do not lead, as they do not for the "traditional" optimal income tax literature, to firm policy recommendations. They do, however, suggest the nature of optimal arrangement in some fairly general classes of circumstances.

Our positive results can be briefly summarized. If total benefits are independent of the income distribution and relative benefits are determined by before- or after-tax income, one should select those projects that yield the greatest total of unweighted benefits across the population. If benefits depend on the distribution of income, the optimal program will be one which produces maximal net benefits at the income distribution which is being induced. Redistribution is a concern, but is carried out solely through the tax system.

References

Atkinson, A. B.: How progressive should income tax be? Chapter 6 in M. Parkin and A. R. Nobay (eds.), *Essays in modern economics*, 1973.

Fair, R. C.: The optimal distribution of income. *Quarterly Journal of Economics 85*, 551–579, 1971.

Feldstein, M.: On the optimal progressivity of the income tax. *Journal of Public Economics 2*, 357–376, 1973.

Harsanyi, J. C.: *Rational behavior and bargaining equilibrium in games and social situations*. Cambridge University Press, 1977.

Mirrlees, J. A.: An exploration in the theory of optimum income taxation. *Review of Economic Studies 38*, 175–208, 1971.

Musgrave, R. A.: *The theory of public finance*. McGraw-Hill, New York, 1959.

Zeckhauser, R.: Uncertainty and the need for collective action. In *The analysis and evaluation of public expenditures: The PPB system*. Joint Economic Committee, U.S. Congress, 1969. Reprinted as Chapter 4 in R. Haveman and J. Margolis (eds.), *Public expenditure and policy analysis*. Markham, Chicago, 1970.

Zeckhauser, R.: Risk spreading and distribution. In H. M. Hochman and G. E. Peterson (eds.), *Redistribution through public choice*. Columbia University Press, New York, 1974.