

Implementation of First-Best Allocations via Generalized Tax Schedules*

THOMAS PIKETTY

*Ecole Normale Supérieure - DELTA,
45, rue d'Ulm, 75005 Paris, France*

Received December 31, 1991; revised September 28, 1992

The classical problem of optimal taxation introduced by J. Mirrlees (*Rev. Econ. Stud.* 38 (1971), 175–208) is addressed, and it is argued that by using generalized tax schedules (i.e., schedules which depend not only on individual income but also on the whole profile of incomes announced by the society) a planner can implement any Pareto optimum, which stands in great contrast with the set of second-best Pareto optima attainable via classical tax schedules. The mechanism relies crucially on the fixed probability distribution of characteristics, and on a finite number of agents. *Journal of Economic Literature* Classification Numbers: D78, D82, H21.

© 1993 Academic Press, Inc.

1. INTRODUCTION

The subject of this paper is the theory of optimal taxation (also referred to by Mirrlees [14] as the theory of optimal public policy), and especially the theory of optimal income taxation.

In the absence of non-convexities, the second welfare theorem states that any Pareto optimum can be seen as the outcome of a competitive equilibrium once adequate lump-sum transfers have been made. Unfortunately, a planner needs to know every characteristic of every agent in the society he faces in order to distribute correctly these lump-sum transfers;¹ consequently, the second welfare theorem is not very useful in a world where the planners have a limited knowledge of the individuals, and questions such as “is it possible to reach Pareto optima different from the competitive allocation” or more generally “what can a planner do” remain unsolved.

* This paper was written in the Autumn of 1991 when I was visiting the London School of Economics and Political Science. I thank the LSE for the nice atmosphere, and especially Kevin Roberts, Max Steuer, John Sutton, David Webb, and Alison Hole. I am grateful to Roger Guesnerie and a referee for their comments.

¹ Even if the planner possessed this information, such a public policy would hardly be acceptable for ethical reasons because of its strong lack of anonymity.

The major advance in this area is due to Mirrlees [13], who recognized the incentive effects associated with taxation and thus the trade-off between equity and efficiency considerations that a planner has to take into account while redistributing wealth among agents by means of a tax system. More precisely, it is now well known that the set of allocations attainable via a general income tax schedule (i.e., a function that associates an after tax income, or consumption level, to a before tax income, or simply income) does not include the Pareto optima of the economy (at least not all of them); the Pareto optima among the allocations of this set are usually referred to as second-best Pareto optima, as opposed to first-best Pareto optima attainable via lump-sum transfers. A comprehensive survey of these results can be found in Stiglitz [16]. This is essentially a negative result (equity is in general not compatible with Pareto-efficiency), and this provides an argument in favour of a very limited redistribution. The first question that one needs to address is the robustness of this result, i.e., is it really impossible for a planner using more sophisticated mechanisms than classical tax schedules to do better than second-best allocations?

A first answer to this question was proposed by Hammond [8] and Guesnerie [6]: with a continuum of agents, and if the characteristics of each agent are drawn at random and independently from the same initial distribution, then a planner cannot do better than using classical tax schedules. A counterpart of this result in large finite economies, based on the same independence assumption, has been proposed by Dierker and Haller [5]. This principle is sometimes referred to as the taxation principle (see Guesnerie [7]). Our argument in this paper is that the independence assumption of these results is too restrictive as compared to what planners can do in the real world, and that, if we adopt a more justified modelling of the information structure of this problem, then a planner can implement much better than second-best allocations by using slightly more complicated mechanisms than classical tax schedules.

First, note that if one wants to design a mechanism which enables planners to go beyond second-best allocations, then one should have a look at another well-alive part of incentive theory which tries to analyze, in very general terms, the set of social choice rules that a planner with a limited knowledge of individual characteristics can implement: the theory of implementation. This theory proved that unfortunately very few social choice rules can be implemented with dominant strategy equilibrium as the equilibrium concept (see Dasgupta *et al.* [4] among others), and in the last decade (especially the most recent years) an important literature established that by using weaker equilibrium concepts (Nash equilibrium, bayesian Nash equilibrium,...) one could go beyond this negative result (see Maskin [11] and Jackson [10]), and even implement almost anything (see Palfrey and Srivastava [15], Abreu and Sen [3] and Abreu and

Matsushima [1, 2]). However, one can easily argue that many mechanisms proposed by this literature lack robustness: the intensive use of integer games (every mechanism à la Maskin), tail-chasing (Palfrey and Srivastava [15]), stochastic outcomes (the essence of virtual implementation) is somewhat confusing; the two former tricks, hardly usable on a large scale, are addressed by Jackson [9], while the credibility of the latter in a general equilibrium setting is questioned by Stiglitz [16, p. 1013]. The size of the message spaces provides another reason why the classical tax system is practically implementable while these mechanisms are not, simply because the transmission of a profile of utility functions might be much less tractable than the transmission of the income (a real number, or more generally a euclidean vector), regardless of the equilibrium robustness. Moreover, an important part of the literature deals with complete information environments, i.e., relies heavily on the assumption that each individual agent knows perfectly the characteristics of everybody in the society while the planner does not; such an assumption is obviously unacceptable in a general equilibrium setting.

The problem we address in this paper is therefore the following: is it possible to design more sophisticated mechanisms (possibly with a weaker equilibrium concept than dominant strategy equilibrium) than non-linear tax systems in order to do better than second-best allocations, while at the same time keeping the practical implementability of the classical taxation? More precisely, we are concerned with the standard two-good income taxation model of Mirrlees [13], with a finite number of agents and a fixed probability distribution of characteristics. Our main argument is that, instead of using a tax schedule (i.e., a function that associates a consumption level to any possible income), a planner should use a schedule that depends not only on the income of the given taxpayer but also on the whole profile of incomes of the society; we will refer to such a schedule as a generalized tax schedule (GTS in the rest of the paper). An agent facing a GTS will not in general have any dominant strategy (as he does not know the incomes that will be announced by the other agents), and thus we will use the equilibrium concept of bayesian Nash equilibrium. Under the single informational assumption that the probability distribution of characteristics is common knowledge, and under the usual Spence-Mirrlees property of indifference curves, we will prove that any first-best allocation can be implemented via GTS; moreover, the (unique) bayesian Nash equilibrium will be formed of the unique iteratively non-strongly dominated strategies, and in fact these will be very simple ones (the common knowledge assumption will be very partially used in the iterative deletion process).

The paper is organized as follows: Section 2 presents the model we are concerned with. Section 3 gives definitions of different implementation

concepts in this framework. Section 4 presents our main result, i.e., the fact that any Pareto-optimal social choice rule can be implemented via generalized tax schedules. Section 5 deals with the assumption of public knowledge of the characteristic distribution, which appeared to be essential for the results of Section 4, and explains why the independence assumption (necessary for the taxation principle to hold) is not a good one in this implementation problem. Section 6 provides concluding comments.

2. THE MODEL

We consider an economy with a set of agents $I = (1; \dots; n)$, two goods C and L , where C is called “consumption” and L is “efficient-labour,” and (for simplicity) a linear technology such that one unit of efficient labour produces one unit of consumption. We consider a finite set of characteristics $\Theta = (\theta_1; \dots; \theta_r)$, and a set of possible individual preferences over (consumption, efficient-labour) bundles, represented by well-behaved (say, continuous, strictly monotonic and strictly quasi-concave) utility functions $(U(C, L, \theta_s), 1 \leq s \leq r)$ parametrized by characteristics. Without any loss in generality, we assume $0 < \theta_1 < \dots < \theta_r$. Each agent $i \in I$ has a characteristic $\theta(i) \in \Theta$, and we denote by $\mu = (\mu(\theta_1); \dots; \mu(\theta_r))$ the distribution of characteristics over the set of agents, where $\mu(\theta_s)$ is the number of agents whose characteristic is θ_s . Without any loss in generality, we assume that $\mu(\theta_s)$ is strictly positive for any $s \in (1; \dots; r)$. The set of possible preferences profiles is

$$P = \{(\theta(i), 1 \leq i \leq n) \text{ such that } \#(i \in I \text{ s.t. } \theta(i) = \theta_s) = \mu(\theta_s) \\ \text{for any } s \in (1; \dots; r)\}.$$

In the income taxation framework, the usual interpretation of individual characteristics is in terms of individual productivities (see Mirrlees [13]): it is usually assumed that there exists a well-behaved utility function $U(C, 1)$ such that for any $\theta \in \Theta$ we have

$$U(C, L, \theta) = U(C, L/\theta).$$

One can think of $1 = L/\theta$ as being an objective measure of disutility due to labour (for instance labour time or effort) appreciated by every agent in the same manner, and of the θ 's as being the abilities or the productivities as they are priced by the private sector (for instance θ can be the wage rate). In this interpretation, the limited knowledge of the planner means that he can only observe the income L (“efficient-labour”) but not the labour time or level of effort l (“effective-labour”). In such a setting, a social planner

might well be interested in redistributing wealth from high ability agents to low ability agents.

An interesting and very useful property of this family of utility functions is that the marginal rates of substitution between consumption and efficient-labour are for every bundle a decreasing function of ability:

$$\forall (C, L), \quad \partial_{\theta}(U_C(C, L, \theta)/U_L(C, L, \theta)) < 0.$$

This is the so-called Spence–Mirrlees condition of principal–agent theory, and it is verified here at least as long as “well-behaved” for $U(C, 1)$ means $U_C > 0$, $U_L < 0$, $U_{CC} < 0$, $U_{LL} < 0$, $U_{CL} < 0$ (note that $U_{CL} < 0$ is not necessary; see Lemma 1, Section 4).

This particular specification of the utility functions is the standard one; even if it is intuitively appealing, this is still a rather restrictive assumption. It can be justified by remarking that if the planner does not make any assumptions of this kind about the inequalities between individuals, then he will hardly be able to make any normative judgement on the different allocations. Anyway, although in the rest of the paper we will refer to this particular specification of the model (in particular we will sometimes refer to “effective-labour” $l = L/\theta$ as well as to “efficient-labour” $L = \theta l$ to represent allocations), it is worth noting that the results of the paper can be readily applied to any family of utility functions $(U(C, L, \theta_s), 1 \leq s \leq r)$ satisfying the Spence–Mirrlees property stated above, whatever the economic interpretation.

The set of possible allocations is defined by

$$A = \left(((C(i), L(i) = \theta(i) l(i)), 1 \leq i \leq n) \right. \\ \left. \text{such that } \sum_{i \in I} C(i) \leq \sum_{i \in I} L(i) \right).$$

In this very simple economy, the “competitive equilibrium” in the absence of government intervention $((C(i), L(i) = \theta(i) l(i)), 1 \leq i \leq n)$ is defined by

$$\forall i \in I, (C(i), L(i)) = \text{ArgMax } U(C, L, \theta(i)) \quad \text{under } C \leq L, \\ \text{or } (C(i), l(i)) = \text{ArgMax } U(C, l) \quad \text{under } C \leq \theta(i)l.$$

(Note that this “competitive equilibrium” is simply the outcome of individual maximizations without any market coordinating any trading).

This competitive allocation is Pareto-optimal, by the first welfare theorem, but this is not the only one. For any $p \in P$, let us define $A(p)$ as the set of Pareto optima of the economy. The second welfare theorem

states that for any $p = (\theta(i), i \in I) \in P$, for any $a = ((C(i), L(i)), i \in I) \in A(p)$, there exists a vector of transfers $(T(i), i \in I) \in \mathbb{R}^n$ such that

- (a) $\sum_{i \in I} T(i) = 0$.
- (b) $\forall i \in I, (C(i), L(i)) = \text{ArgMax } U(C, L, \theta(i))$ under $C \leq L + T(i)$.

If the planner knows μ but cannot observe p , then he is unable to distribute the transfers correctly. What can he do if he is however not satisfied with the competitive allocation implemented for $T(i) = 0 \forall i \in I$ (for instance for equity reasons because he considers that individuals are not responsible for their low wage rate)? Before trying to answer to this question, a few other definitions are necessary.

The normative judgements of the planner are embodied in a social choice rule Ψ , which is a mapping from P to A . We say that a social choice rule is anonymous if it depends only on the characteristics of the agents, not on their "name":

$$(\Psi: P \rightarrow A \text{ anonymous}) \\ \Leftrightarrow (\forall p = (\theta(i), i \in I), p' = (\theta'(i), i \in I) \in P, \forall i, j \text{ s.t. } \theta(i) = \theta'(j),$$

then

$$(C(i), L(i)) = (C'(j), L'(j)),$$

where

$$\Psi(p) = ((C(k), L(k)), k \in I)$$

and

$$\Psi(p') = ((C'(k), L'(k)), k \in I).$$

Note that an anonymous social choice rule is thus completely determined by a vector $((C_s, L_s), 1 \leq s \leq r)$, where (C_s, L_s) is the bundle for θ_s -agents (i.e., the agents whose characteristic is θ_s). In what follows we will use both formulations and refer to the latter as an "anonymous allocation."

3. THE IMPLEMENTATION PROBLEM IN THIS FRAMEWORK

To the question "what can a planner do in such an environment", implementation theory answers that the most general thing he can do is to design a game form (alternatively called a "mechanism"), i.e., a strategy space S_i for each agent and a mapping π from the vectors of strategies to A , such that for any preference profile the individually rational strategies

lead to the outcome chosen by the planner for that particular preference profile. Thus the problem depends crucially on what one considers to be an individually rational choice of strategies, i.e., on the equilibrium concept used. If one chooses dominant strategy equilibrium as the norm of rational behavior, the definition of implementation is the following:

Given a social choice rule Ψ , a game form (S, π) is said to implement Ψ in dominant strategy equilibrium if and only if:

$$\forall p = (\theta(i), i \in I) \in P, \forall i \in I, \exists s_i(p) \in S_i \text{ such that}$$

$$(a) \quad \forall s_i \in S_i, s_{-i} \in S_{-i}, U(\pi(s_i(p), s_{-i}), \theta(i)) \geq U(\pi(s_i, s_{-i}), \theta(i))$$

$$(b) \quad \pi(s_i(p), i \in I) = \Psi(p)$$

(c) $\forall (s'_i(p), i \in I)$ verifying (a), $\pi(s'_i(p), i \in I) = \Psi(p)$ (with obvious notations).

In order to define implementation in bayesian Nash equilibrium, one needs to specify completely the informational structure of the model (i.e., what the agents know about the other agents' characteristics, about the other agents' beliefs over these characteristics, and so on...). The weakest and most natural assumption that one can make in this model is to assume that each agent knows only his own type and the distribution μ , and that this is common knowledge. In this context a mechanism (S, π) is said to implement a social choice rule Ψ in bayesian Nash equilibrium if and only if:

$$\forall i \in I, \forall \theta \in \Theta, \exists s_i(\theta) \in S_i \text{ such that}$$

$$(a) \quad \forall s_i \in S_i, \forall \theta \in \Theta, \sum_{\theta(-i)} U(\pi(s_i(\theta), s_{-i}(\theta(-i))), \theta) v(\mu, \theta, \theta(-i)) \geq \sum_{\theta(-i)} U(\pi(s_i, s_{-i}(\theta(-i))), \theta) v(\mu, \theta, \theta(-i))$$

$$(b) \quad \forall p = (\theta(i), i \in I) \in P, \pi(s_i(\theta(i)), i \in I) = \Psi(p)$$

(c) $\forall (s'_i(\theta), i \in I)$ verifying (a), $\forall p = (\theta(i), i \in I) \in P, \pi(s'_i(\theta(i)), i \in I) = \Psi(p)$ (where $v(\mu, \theta, \theta(-i))$ is the distribution over the other agents' characteristics $\theta(-i)$ deduced from the knowledge of μ and θ).

These two definitions define the set of allocations that a planner can implement by using these equilibrium concepts. The aim of implementation theory is to characterize more precisely these sets. However, in this model, where there is much more structure than in the classical implementation framework, there exist more usual and seemingly simpler mechanisms to implement social choice rules: taxation systems. The next and natural problem is then to determine what can be achieved via tax systems, and what their link with general mechanisms is.

A taxation system, from a general mechanism design viewpoint, is characterized by the property that it is strongly anonymous, in the sense

that the (consumption, efficient-labour) bundle allocated to an agent depends only on his strategy. The formal definition of a tax schedule is the following:

A tax schedule is a real function Φ such that:

- (a) $\forall \theta \in \Theta, \exists! (C(\theta), L(\theta)) = \text{ArgMax } U(C, L, \theta)$ under $C \leq \Phi(L)$
- (b) $\sum_{\theta \in \Theta} C(\theta) \mu(\theta) \leq \sum_{\theta \in \Theta} L(\theta) \mu(\theta)$ (L is the before tax income, $\Phi(L)$ is the after tax income).

Note that implementation in dominant strategy equilibrium and implementation in bayesian Nash equilibrium are equivalent as long as one uses strongly anonymous mechanisms. It is straightforward to show that tax schedules are the most general strongly anonymous mechanism: any allocation which can be implemented via a strongly anonymous mechanism can be implemented by using a well-chosen tax schedule.

The taxation principle wants to say more: it states that this restriction to strongly anonymous mechanisms is in fact not a restriction since a planner cannot implement more allocations by using more mechanisms. The taxation principle is presented in more general contexts than the Mirrlees' model by Hammond [8] and Guesnerie [6], among others. The Pareto optima among the set of allocations implementable via tax schedules are called the second-best Pareto optima. In general, they are different from the first-best Pareto optima. A complete description of this set is possible when the utility functions have the particular form considered in Section 2 (see Stiglitz [16]), and constitutes one of the most classical applications of the principal-agent model with adverse selection. It is worth noting that the knowledge of μ by the planner is in fact not necessary in order to implement second-best allocations: for any given set of possible distributions, a planner who does not know which distribution is the right one can implement in dominant strategy equilibrium exactly the allocations implementable when he knows the distribution (see Guesnerie [6]).

The objective of the taxation principle is clear: one should concentrate on the study of tax schedules, since nothing more can be achieved by using more sophisticated mechanisms. However, the taxation principle holds only under the assumption that there is a continuum of agents and that the characteristics of each agent are drawn at random and independantly from the same initial distribution. In the case of finite economies, the taxation principle simply tells us that tax schedules are the most general strongly anonymous mechanism, without justifying at all the restriction to strongly anonymous mechanisms. The objective of this paper is precisely to prove that the strong anonymity restriction (or equivalently the tax schedules restriction) is no longer justified in the finite economy considered here; some more sophisticated tax systems should be designed for (Pareto-)

efficiency reasons. This crucial difference between economies with a continuum of agents and economies with a finite set of agents is discussed in Section 5. Moreover, note that if one wants to limit the set of usable mechanisms because of anonymity considerations, then the correct restriction should be to use only weakly anonymous mechanisms (i.e., mechanisms such that the bundle given to an agent depends only on his strategy and on the probability distribution induced by the other agents' strategies; roughly speaking, the "names" do not matter (the formal definition is left to the reader)). There exist several reasons for doing so: firstly, weak anonymity is the adequate concept as far as political or ethical considerations are concerned, while strong anonymity is needlessly restrictive; secondly, since the planner has exactly the same information about each individual, no Pareto improvement can be obtained by using mechanisms which are not weakly anonymous.²

We define generalized tax schedules as the natural extension of tax schedules derived from weakly anonymous mechanisms exactly in the same way as tax schedules are derived from strongly anonymous mechanisms. More precisely:

A generalised tax schedule (GTS) is a mapping Φ that associates a consumption level to an individual income and a distribution of the other agents' incomes such that:

- (a) $\forall \theta \in \Theta, i \in I, \exists! (C(i, \theta), L(i, \theta)) = \text{ArgMax } U(C, L, \theta) \text{ under } C \leq \Phi(L, v((L(j, \theta'), j \neq i, \theta' \in \Theta), \theta, \mu))$
- (b) $\forall p = (\theta(i), i \in I) \in P \sum_{i \in I} C(i, \theta(i)) \mu(\theta(i)) \leq \sum_{i \in I} L(i, \theta(i)) \mu(\theta(i))$ (where $v(\cdot)$ is the distribution of incomes of the rest of the society induced by the strategies of the other agents).

The definition of the implementation of a social choice rule via a generalized tax schedule follows directly from the definition given above. In particular, it should be clear that implementability via GTS is a special case of implementability in bayesian Nash equilibrium. We can now prove that the move from classical tax schedules to GTS strongly increases the planner's possibilities.

² Following Guesnerie [6], one can distinguish between two types of anonymity: horizontal anonymity and anonymity in influence. Using mechanisms that are not horizontally anonymous is equivalent to *ex ante* randomization over implementable social choices rules, and thus a planner cannot generate any Pareto improvement by doing so, even if it can be interesting for a planner with an utilitarian welfare function, because of potential non-convexities (see Stigitz [16, p. 1012]). Using mechanisms that are not anonymous in influence is equivalent to *ex post* randomization, which does not seem to be interesting, except if one makes additional assumptions on the different types' risk aversion. Consequently, non-weakly anonymous mechanisms, while implementing much more social choice rules, are not interesting as far as Pareto improvements are concerned.

4. FIRST-BEST ALLOCATIONS ARE IMPLEMENTABLE VIA GTS

Before we state and prove our main result, we would like to show with a single case how it works, since the intuition behind the result is very simple.

Let us consider the special case where $r=2$, and $\theta_2 > \theta_1$. Let us consider a first-best anonymous allocation $((C_1, L_1), (C_2, L_2))$ that involves a large wealth redistribution from the θ_2 -agents to the θ_1 -agents so that the θ_2 -agents prefer the bundle (C_1, L_1) initially devoted to θ_1 -agents to their own bundle (C_2, L_2) (see Fig. 1). Typically, this allocation is not implementable via a tax schedule. If a planner using tax schedules really wants to give to the θ_1 -agents the utility $U(C_1, L_1, \theta_1)$ then he will have to lower the utility level of θ_2 -agents and to use a distortionary taxation with θ_1 -agents. Consequently this second-best Pareto optimum (noted $((C_{1SB}, L_{1SB}), (C_{2SB}, L_{2SB}))$ in Fig. 1) will not be a first-best Pareto optimum.

Let now see why a planner using GTS can easily implement the allocation $((C_1, L_1), (C_2, L_2))$. Let C'_1 and L'_1 be the consumption levels such that

$$U(C'_1, L_1, \theta_2) = U(C_2, L_2, \theta_2), \quad U(C'_1, L_1, \theta_1) = U(C_2, L_2, \theta_1).$$




FIGURE 1

Because of the Spence–Mirrlees property, we have $C_1'' < C_1'$. Suppose that the planner say to the agents: if you earn L_2 , you will consume C_2 ; if you earn L_1 and the number of agents whose income equals L_1 is strictly superior to $\mu(\theta_1)$, then you will consume C^* (with $C_1'' < C^* < C_1'$); if you earn L_1 and the number of agents whose income equals L_1 is inferior or equal to $\mu(\theta_1)$, then you will consume C_1 ; if you earn anything different from L_1 or L_2 , then you will consume, say, 0 (or anything preferred by (C^*, L_1) or (C_2, L_2)). Then any θ_1 -agent facing this GTS has a strictly dominant strategy: earning L_1 . Anticipating that the θ_1 -agents will earn L_1 , any θ_2 -agent has a unique best-response (whatever the strategies of the other θ_2 -agents): earning L_2 . Consequently, this GTS implements the first-best anonymous allocation $((C_1, L_1), (C_2, L_2))$.

The implementation of any anonymous social choice rule $((C_l, L_l), 1 \leq l \leq r)$ via GTS works exactly in the same way: we design a GTS such that (a) the type 1-agents have a strictly dominant strategy (earning the income planned for their type); (b) the 2-agents, knowing that the first type-agents will play their dominant strategy, have only one strategy which is not strongly dominated (earning the income planned for their type); (c) and so on... .

Being formed of the unique strategies which survive the iterative removal of strongly dominated strategies, this equilibrium is the unique bayesian Nash equilibrium: to prove this general game theory result in the special case of our very simple iterative deletion process, note firstly that any bayesian equilibrium must involve the strictly dominant strategy of type 1-agents (since by definition this strategy is strictly better than any other strategy, whatever the other agents' strategies); secondly any bayesian equilibrium must also involve the strategy of type 2-agents which strictly dominates any other strategy once the other agents' strategies vector involve the strictly dominant strategy of type 1-agents (since we already know that this is the case for the other agents' strategies in any bayesian equilibrium); and so on... .

We consider this equilibrium structure as a crucial argument in favour of GTS, especially if one compares it to the confusing and complicated tricks often used to rule out unwanted equilibria. Moreover, as one can check, only a very partial part of the common knowledge assumption is used in the iterative deletion process, and the planner knows which agents' type needs to have some true beliefs of a superior order. From a more conventional game theoretic viewpoint, note that this refinement of Nash equilibrium is the most robust one (it is perfect in any sense of this term); moreover it is well known that it is implied by bayesian rationality (i.e., utility maximization given an assessment of other players' strategies) and common knowledge of bayesian rationality, so that the removal of strictly dominated strategies should be

considered as the rationality concept coming just after dominant strategy rationality.

As can be intuitively understood by looking at Fig. 1, a necessary condition for designing such a mechanism is that the before tax income L_2 of the high productivity agents is larger than the before tax income L_1 of the low productivity agents. However, this is a very weak restriction: the Spence–Mirrlees condition implies that the competitive allocation in the absence of government intervention involves a higher before tax income for higher productivity agents, so that if we assume leisure to be a non-inferior good we get that every (first-best) Pareto optimum characterized by higher transfers for lower productivity agents involves higher before tax incomes for higher productivity agents. Since there is a little interest from a normative point of view for Pareto optima characterized by higher transfers for higher productivity agents, virtually every Pareto optimum (say, every acceptable Pareto optimal social choice rule) can be implemented with our mechanism, even if strictly speaking Pareto optima characterized by a sufficiently large transfer from low productivity agents to high productivity agents may well involve higher before tax income for lower productivity agents. Anyway, regardless of this discussion, we have, as mentioned above:

LEMMA 1. *Let $((C_s, L_s) = \text{ArgMax } U(C, L, \theta_s)$ under $C \leq L + T_s$, $1 \leq s \leq r$) be the Pareto-optimal anonymous allocation generated by a lump-sum transfer vector $(T_s, 1 \leq s \leq r)$ such that $\sum_{1 \leq s \leq r} T_s = 0$ and $T_s \geq T_t$ for $s \leq t$. Let us assume that the Spence–Mirrlees condition holds and that $U(C, l)$ is such that leisure (i.e., $-l$) is a non-inferior good. Then $L_s > L_t$ for $s > t$.*

Proof. The Spence–Mirrlees property implies that this monotonicity property holds for an anonymous allocation $((C_s, L_s) = \text{ArgMax } U(C, L, \theta_s)$ under $C \leq L + T$, $1 \leq s \leq r$) for any real number T (see Maskin and Riley [12, Assumption 3, pp. 4–5]). The very definition of non-inferiority implies that if $((C, L) = \text{ArgMax } U(C, L, \theta)$ under $C \leq L + T$) and $((C', L') = \text{ArgMax } U(C, L, \theta)$ under $C \leq L + T'$), with $T > T'$, then $L < L'$. This completes the proof. Note also that the Spence–Mirrlees property and the non-inferiority of leisure are both implied by the usual assumptions $U_C > 0$, $U_l < 0$, $U_{CC} < 0$, $U_{ll} < 0$, $U_{Cl} < 0$ (these assumptions are however not necessary; the sufficient and necessary conditions on $U(C, l)$ for these two properties are given by Maskin and Riley [12, pp. 3–5]). C.Q.F.D.

We can now state and prove our main result:

PROPOSITION 1. *Let $((C_s, L_s), 1 \leq s \leq r)$ be an anonymous allocation such that $L_s > L_t$ for $s > t$. Then there exists a GTS Φ which implements it via*

iterative removal of strongly dominated strategies (and therefore in bayesian Nash equilibrium)

Proof. Let $((C_s, L_s), 1 \leq s \leq r)$ be an anonymous social choice rule. Let us assume that there exists a sequence of consumption levels $(C_{st}, 1 \leq s, t \leq r)$ such that:

$$\begin{aligned} \forall s, t, u \text{ such that } 1 \leq s \leq r, s \leq u \leq r, 1 \leq t \leq r, t \neq s, \text{ then} \\ C_{ss} = C_s \tag{1} \\ U(C_{su}, L_s, \theta_s) > U(C_{tu}, L_t, \theta_s). \end{aligned}$$

Let us define a generalized tax schedule Φ by

- (a) $\Phi(L, v)C_{st}$ if $L = L_s$ and
 $t = \text{Min}\{t' \text{ s.t. } 1 \leq t' \leq r \text{ and } v(L_{t'}) < \mu(\theta_{t'})\}$,
for any $1 \leq s, t \leq r$.
- (b) $\Phi(L, v) = 0$ if $L \neq L_s$ for any $s \in (1; \dots; r)$.

Note that this schedule is well defined since for any possible distribution v of the other agents' incomes there exists a unique $t = \text{Min}\{t' \text{ s.t. } 1 \leq t' \leq r \text{ and } v(L_{t'}) < \mu(\theta_{t'})\}$.

Consider an agent i such that $\theta(i) = \theta_1$. Then, using (1), we have

$$\forall L \geq 0, \quad L \neq L_1, \quad \forall v, \quad U(L_1, \Phi(L_1, v), \theta_1) > U(L, \Phi(L, v), \theta_1).$$

Therefore L_1 is a strictly dominant strategy for any θ_1 -agent. Let us now consider an agent j such that $\theta(j) = \theta_2$. Because of the assumption that j knows $\mu(\theta_1)$, j knows that he will never face any income distribution v such that $v(L_1) < \mu(\theta_1)$. Moreover, using (1),

$$\begin{aligned} \forall L \geq 0, \quad L \neq L_2, \quad \forall v \text{ s.t. } v(L_1) \geq \mu(\theta_1), \\ U(L_2, \Phi(L_2, v), \theta_2) > U(L, \Phi(L, v), \theta_2). \end{aligned}$$

Thus L_2 is the unique iteratively non-strongly dominated strategy for any θ_2 -agent. Repeating these remarks until $\theta(k) = \theta_r$ leads to the fact that L_s is the unique strategy which survives the iterative removal of strongly dominated strategies for any θ_s -agent, for any $s \in (1; \dots; r)$. It follows that if we prove the existence of a sequence $(C_{st}, 1 \leq s, t \leq r)$ verifying (1), then the proof of the Proposition 2 is complete.

Consider the following proposition, depending on an integer $n \in (1; \dots; r)$:

$$P(n) : \forall ((\theta_s^*, L_s^*), 1 \leq s \leq n), C^* > 0 \text{ such that } \theta_s^* < \theta_t^* \Leftrightarrow s < t \quad (2)$$

$$L_s^* < L_t^* \Leftrightarrow s < t \quad (3)$$

There exist $(C_s^{**}, 1 \leq s \leq n)((\theta_s^*, L_s^*), 1 \leq s \leq n), C^*)$ such that:

$$C_n^{**} = C^* \quad \forall s, t \text{ such that } 1 \leq s, t \leq n, t \neq s, \text{ then}$$

$$U(C_s^{**}, L_s^*, \theta_s^*) > U(C_t^{**}, L_t^*, \theta_s^*).$$

If we prove that $P(n)$ is true for any $n \in (1; \dots; r)$, then we will have proved the existence of the sequence $(C_{st}, 1 \leq s, t \leq r)$ verifying (1) for any anonymous allocation $((C_s, L_s), 1 \leq s \leq r)$ such that L_s is an increasing function of ability: to see this, just consider the $(C_{st}, 1 \leq s, t \leq r)$ defined by

$$\forall t \in (1; \dots; r), \quad (C_{st}, 1 \leq s \leq t) = (C_s^{**}, 1 \leq s \leq t)((\theta_s, L_s), 1 \leq s \leq t), C_t)$$

$$C_{st} = 0 \quad \text{if } s > t.$$

We prove that $P(n)$ is true by induction.

$P(1)$ is trivially true. Assume that $P(n)$ is true. Let us consider any $((\theta_s^*, L_s^*), 1 \leq s \leq n+1), C^*)$ verifying (2) and (3). We can apply $P(n)$ to $((\theta_s^*, L_s^*), 2 \leq s \leq n+1), C^*)$. For any s such that $2 \leq s \leq n+1$, consider:

$$C'_s, C''_s \text{ s.t. } U(C'_s, L_1^*, \theta_s^*) = U(C_s^{**}, L_s^*, \theta_s^*)$$

$$U(C''_s, L_1^*, \theta_1^*) = U(C_s^{**}, L_s^*, \theta_1^*)^3$$

(where the C_s^{**} stand for the $(C_s^{**}, 2 \leq s \leq n+1)((\theta_s^*, L_s^*), 2 \leq s \leq n+1), C^*)$ obtained by the application of $P(n)$).

We claim that $\max(C''_s, 2 \leq s \leq n+1) < \min(C'_s, 2 \leq s \leq n+1)$.

To prove this, suppose the opposite, i.e., suppose that there exists s, t such that $C'_s < C''_t$.

We can now prove that this leads to a contradiction with the Spence–Mirrlees property (see Fig. 2): consider the real mappings $F(L)$ and $G(L)$ defined by

$$\forall L \in [L_1^*, L_t^*] \quad U(G(L), L, \theta_s^*) = U(C_s^{**}, L_s^*, \theta_s^*)$$

$$U(F(L), L, \theta_1^*) = U(C_t^{**}, L_t^*, \theta_1^*)$$

³ C'_s and C''_s are well defined, as a consequence of the strict monotonicity of preferences (so that then any indifference curve defines a strictly increasing real mapping).




FIGURE 2

(these mappings are continuous and strictly monotonic due to continuity and strict monotonicity of preferences) $P(n)$ gives us:

$$U(C_s^{**}, L_s^*, \theta_s^*) > U(C_t^{**}, L_t^*, \theta_s^*)$$

Therefore $C_t^{**} = F(L_t^*) < G(L_t^*)$. $C_s < C_t'$ implies $F(L_1^*) > G(L_1^*)$. The continuity of the mappings $F(L)$ and $G(L)$ implies the existence of $L' \in [L_1^*, L_t^*]$ such that (a) $F(L') = G(L')$ and (b) $F'(L') \leq G'(L')$. But this contradicts the Spence–Mirrlees condition at the point $(F(L'), L')$, since $\theta_1^* < \theta_s^*$.

Thus we have proved the existence of a C_1^{**} such that

$$\forall s \in (2; \dots; n+1), C_s'' < C_1^{**} < C_s''.$$

Now, one can check that the $(C_s^{**}, 1 \leq s \leq n+1)$ satisfy the definition of the $(C_s^{**}, 1 \leq s \leq n+1)((\theta_s^*, L_s^*), 1 \leq s \leq n+1), C^*$.

Thus we have proved $\bar{P}(n+1)$.

C.Q.F.D.

It is worth noting that it is possible to smooth the GTS defined in the proof, exactly as in the case of tax schedules. In particular, the (b) of the definition of the GTS was only used for convenience and can be removed (see the example with two characteristics).

5. ON THE INFORMATION ASSUMPTION

While the common knowledge assumption is not fully necessary (see the remarks just before Lemma 1), it appeared clearly that the public knowledge of the characteristics distribution is essential for the efficiency of the GTS. One can wonder what the meaning of this assumption is, and whether this is a reasonable one.

Concerning the knowledge of the true distribution by the planner, two simple justifications can be provided.

First, it can be argued that any planner who wants to do some wealth redistribution inside a given society needs to have (and has in the real world) a statistical knowledge of this society. Of course, this knowledge is very unlikely to be a perfectly precise one, while our mechanism seems to rely heavily on the assumption that the planner knows the exact μ in order to define properly his GTS. In fact, it does not, as the planner can always design a GTS using the superior bound of possible sizes of the set of agents whose allocations is envied; by doing so, the planner will implement an allocation only slightly different from the Pareto optimum (in the example with two characteristics, there might be a few θ_2 -agents who will safely deviate) as long as the planner's estimation of the true distribution is not too bad. This "slight difference" should not be so important for a planner facing a large economy. These allocations obtained via "GTS with a small margin of error" should anyway be much better than second-best allocations. This "continuity" intuition must of course be made more precise.

Second, a planner without any prior on the characteristics distribution should manage to solve this problem in a dynamic sense (and the real world planners who do not have sufficient information certainly do so). To solve theoretically such a problem, the fact that the implementation in dominant strategy equilibrium of second-best allocations does not need the knowledge of the distribution (see Section 3) might be useful: one could think to a two-stage mechanism where in the first period a second-best allocation is implemented in dominant strategy equilibrium and a second period where a first-best allocation is implemented via GTS by using the information collected during the first period. Of course, this has also to be made more precise.

Concerning the knowledge of the distribution by the agents, this should not be a real problem as long as the planner has the information: he can simply communicate it to the agents, as long as it is in their own interest to believe him. Even if difficult related problems are still to be solved (for example if the agents' "small error" in their assessment of the true distribution is in a different direction than the planner's error), the argument of the remark above concerning the small deviations due to an imprecise prior should simplify some of them.

Another point is worth discussing: it is somewhat confusing that the knowledge of μ is no longer usable by GTS when one models the set of agents as a continuum (see Section 3 and Guesnerie [6]).⁴ The reason is that the continuum assumption rules out the possibility for a planner of allocating in a different way after an individual deviation, while this possibility exists in the real world (even if a planner can only use it with a margin of error, as noted above, this threat is enough to deter (most of) the agents from choosing a “wrong” income in equilibrium). Indeed, assuming that individual characteristics are drawn from independent random variables implies for example that the observation of the characteristics of an important part of the society can never improve our statistical knowledge of the other part of the society; however, in the real world, if a planner faces for example a society almost equally divided between “rich” and “poor”, and if he observes that a first half of the society is, say, poor, then this will greatly improve his statistical knowledge of the second half of the society, even if he is not completely sure that everybody in this second half is “rich”; this possibility for improving information, which is ruled out by the independence assumption, exists in the real world, should be allowed by the modelling, and is extensively used by our mechanism. For these reasons, we do think that models with a finite number of agents and a fixed distribution of characteristics are more justified ones for this problem.

Note that these remarks can be made from a more general implementation theory viewpoint: the continuum assumption leads to the extreme case where the information is completely private, so that the usual incentive compatibility constraints are fully binding; whereas in a finite economy, we are just in the opposite case: the information is non-exclusive (that is anyone’s type can be inferred from the collective information of the other agents), and this is why much more can be implemented.

Note that the information would still be non-exclusive for example if the planner did not know the distribution of types but instead (several of) the agents knew it, and hence incentive compatibility constraints would still be avoided: in fact, Abreu and Matsushima’s [3] mechanism implements any social choice rule if the information is non-exclusive, but this relies on the concept of virtual implementation (see Section 1), and in the context of non-virtual implementation non-exclusivity of information does not imply that anything can be implemented (see Jackson [10]). Anyway, regardless of the fact that this information assumption does not seem very suitable in

⁴ As noted above, Dierker and Haller [5] proposed an asymptotic result which extends the domain of application of the taxation principle to large finite economies; however, the independence assumption behind the result is exactly the same as under the continuum modelization.

the income taxation framework (tax-designers certainly have a better statistical knowledge of the population than the agents), it seems clear that a GTS, which provides a much less sophisticated mechanism than general game-forms, is only suitable for a mechanism-designer who knows the distribution of types (at least approximately): it does not seem to be possible for a planner to extract the distribution of types from the agents by using a GTS, even if it may be feasible with another mechanism.

6. CONCLUDING COMMENTS

The results obtained for implementation via GTS are obviously partial, even if the special type of economy considered in this paper is in fact a very important one as far as the opportunity of redistribution is concerned. Moreover, it seems very unlikely that in the general case the use of GTS does not enable a planner to do much better than with simple tax schedules.

Second-best allocations are supposed to be the best allocations attainable by a planner in a context of information asymmetry. In the light of our results, this definition should be qualified a bit: second-best allocations are the allocations attainable in such a context by a planner who restricts himself to the use of tax schedules. As the very refinement of bayesian Nash equilibrium formed by the equilibrium of our mechanism should be considered as the smallest weakening of dominant strategy rationality, it is unclear whether the restriction which founds the actual understanding of second-best allocations is justified.

REFERENCES

1. D. ABREU AND H. MATSUSHIMA, Virtual implementation in iteratively undominated strategies: Complete information, mimeo, 1990.
2. D. ABREU AND H. MATSUSHIMA, Virtual implementation in iteratively undominated strategies: Incomplete information, mimeo, 1990.
3. D. ABREU AND A. SEN, Virtual implementation in Nash equilibrium, *Econometrica* **59** (1991), 997-1021.
4. P. DASGUPTA, P. HAMMOND, AND E. MASKIN, The implementation of social choice rules: Some general results in incentives compatibility, *Rev. Econ. Stud.* **46** (1979), 185-216.
5. E. DIERKER AND H. HALLER, Tax systems and direct mechanisms in large finite economies, mimeo, Vienna, 1989.
6. R. GUESNERIE, "On Taxation and Incentives: Further Reflections on the Limits to Redistribution," WP Bonn, 1981.
7. R. GUESNERIE, "The Arrow-Debreu Paradigm Faced with Modern Theories of Contracting: a Discussion of Selected Issues Involving Information and Time," Nobel Symposium No. 77, Stockholm, August 1990, and WP DELTA No. 90-26, 1990.

8. P. HAMMOND, Straightforward incentive compatibility in large economies, *Rev. Econ. Stud.* **46** (1979), 263–282.
9. M. JACKSON, *Implementation in undominated strategies: A look at bounded mechanisms*, mimeo, Northwestern University, 1989.
10. M. JACKSON, Bayesian implementation, *Econometrica* **59** (1991), 461–477.
11. E. MASKIN, The theory of implementation in Nash equilibrium, in “Social Goals and Social Organization” (Hurwicz, Schmeidler, and Sonnenschein, Eds.), 1985.
12. E. MASKIN AND J. RILEY, Input versus output incentive shemes, *J. Public. Econ.* **28** (1985), 1–25.
13. J. MIRRLEES, An exploration in the theory of optimal income taxation, *Rev. Econ. Stud.* **38** (1971), 175–208.
14. J. MIRRLEES, The theory of optimal taxation, in “Handbook of Mathematical Economics,” Vol. 3, 1986.
15. T. PALFREY AND S. SRIVASTAVA, Nash implementation using undominated strategies, *Econometrica* **59** (1991), 479–501.
16. J. STIGLITZ, Pareto efficient and optimal taxation and the new and welfare economics, in “Handbook of Public Economics,” Vol. 2, 1987.