INDIRECT TAXATION AND REDISTRIBUTION: THE SCOPE OF THE ATKINSON-STIGLITZ THEOREM

by

Robin Boadway, Queen's University, Canada

and

Pierre Pestieau, Université de Liege, Belgique

1. Introduction

One of the oldest controversies in tax theory involves the choice between direct and indirect taxation, in particular the issue of when differential commodity taxes are not a component of the optimal tax system. The early literature focussed on the efficiency role of commodity taxes: under what circumstances would the Ramsey tax system applied to a given household consist of a uniform tax on commodities, or equivalently a tax on income? The famous Corlett and Hague (1953-54) Theorem settled that. If all goods are 'equally substitutable' for leisure, differential commodity taxes should not be used. Otherwise, goods that are more complementary with leisure should bear higher commodity tax rates. As explained in Sandmo (1976), a utility function in which goods are separable from leisure, and which is homothetic in goods satisfies this property. This result, although an important methodological innovation, is of limited interest from a policy point of view since it abstracts from the redistributive role that the tax system plays.

The question of when differential commodity taxes should be used alongside a progressive income tax as part of a redistributive tax system was addressed in a well-known paper by Atkinson and Stiglitz (1976). Their result, the *Atkinson-Stiglitz Separability Theorem* (the A-S Theorem in what follows), has been seminal and has spawned a substantial literature. Roughly speaking, the A-S Theorem states that if household utility functions are separable in goods and leisure, differential commodity taxes should not be used. This result is arguably the most relevant result for policy purposes to emerge from the optimal income tax literature initiated by Mirrlees (1971). It has been subject to considerable scrutiny in the literature, and special attention has been devoted to the circumstances in which it is violated and what it implies for the structure of commodity taxes.[1] Interestingly, the analogue of the Corlett-Hague Theorem applies, albeit for different reasons. As shown by Edwards et al (1994) and Nava et al (1996), if weak separability is violated, higher tax rates should apply to goods that are relatively more complementary with leisure.

Our purpose in this paper is to revisit the A-S Theorem. We explore the robustness of the theorem to different specifications of household utility, including differences among households in needs and preferences as well as different types of labor supply. We show that the applicability of the A-S Theorem depends crucially on the information available to the government. We begin with a simple derivation of the A-S Theorem, using a methodology that will be useful in synthesizing the various extensions. We then turn to those extensions, first focusing on the case of different preferences and needs for particular goods, and then turning to the case where households can supply alternative forms of labor, including non-market labor.

2. The A-S Theorem

In this section, we adopt a simplified version of the model used by Atkinson and Stiglitz (1976), retaining their essential assumptions. There are two types of households who differ only in their wage rates w_i $(i = 1, 2)$, where $w_2 > w_1$, with n_i households of type i. We assume there are only two goods, denoted x and z, along with labor ℓ, and that households have identical weakly separable utility functions of the form $u(g(x, z), \ell)$.[2] The utility function is strictly concave, and both goods as well as leisure are normal. The market (pre-tax) income of a type-i household is $y_i \equiv w_i \ell_i$. Following Guesnerie (1995), the government is assumed to be able to observe household incomes as well as anonymous transactions in the goods market. It can therefore implement a non-linear income tax as well as proportional commodity taxes.[3] As is well known, only the structure of commodity taxes, and not their level, constitutes an independent policy instrument: proportional commodity taxes can be replicated by an appropriate adjustment in the income tax schedule. Therefore, we can normalize the commodity tax rate on good x to be zero, and treat the tax rate on z as the policy instrument reflecting the differential commodity tax structure. Let t be the per unit tax on purchases of good z. If $t = 0$ in the optimum, the redistributive objectives of government can be achieved by an income tax alone. Goods prices are normalized to unity, and we define the consumer price of good z to be $q \equiv 1 + t$.

To facilitate our analysis, we disaggregate household decision-making into two stages.[4] In the first stage, the household chooses labor supply, earns income, pays income taxes, and ends up with disposable income. In the second stage, disposable income is allocated between the two goods. Consider the second stage first. Let c_i be disposable income, where $c_i = x_i + qz_i$. Given the separable utility function, a household of type i solves the following problem:

$$\max_{\{z_i\}} \quad g\left(c_i - qz_i, z_i\right)$$

where c_i and q are given. From the first-order conditions, $g_z^i / g_x^i = q$,[5] we obtain the demand function $z(q, c_i)$ and the value function $h(q, c_i)$. Applying the envelope theorem with respect to q and c_i, we obtain:

$$h_q^i = -g_x^i z_i, \qquad \text{and} \qquad h_c^i = g_x^i$$

In the first stage, the household chooses labor supply, given the income tax schedule chosen by the government and the anticipated outcome of stage 2. Effectively, the household is choosing earned income y_i and, via the income tax, c_i. For this stage, we follow the standard procedure of optimal income tax analysis initiated by Stiglitz (1982) of allowing the government to choose y_i and c_i implicitly by its choice of an income tax schedule. Individual utility functions are reformulated in terms of what the government can observe as follows:

$$v^i\left(h\left(q,c_i\right),\frac{y_i}{w_i}\right)\equiv u\left(h\left(q,c_i\right),\ell_i\right)$$

The government is assumed to maximize a utilitarian objective function, although any quasi-concave function in individuals utilities would give the same results. The Lagrange expression for the optimal income and commodity tax problem of the government can then be written as:

$$\mathcal{L}=\sum_{i=1,2}n_i v^i\left(h\left(q,c_i\right),\frac{y_i}{w_i}\right)+\lambda\sum_{i=1,2}n_i\left(y_i+tz^i\left(q,c_i\right)-c_i\right)$$

$$+\gamma\left[v^2\left(h\left(q,c_2\right),\frac{y_2}{w_2}\right)-v^2\left(h\left(q,c_1\right),\frac{y_1}{w_2}\right)\right]$$

The first constraint with Lagrange multiplier λ reflects the government budget constraint, and assumes no net revenue requirement. The second constraint with multiplier γ is the incentive constraint and reflects the fact that this will only be binding for type-2 households.

The relevant first-order conditions for our purposes are those with respect to c_1, c_2 and q:[6]

$$n_1v_h^1h_c^1-\lambda n_1\left(1-t\frac{\partial z^1}{\partial c_1}\right)-\gamma\widehat{v}_h^2\widehat{h}_c^2=0 \tag{1}$$

$$n_2v_h^2h_c^2-\lambda n_2\left(1-t\frac{\partial z^2}{\partial c_2}\right)+\gamma v_h^2h_c^2=0 \tag{2}$$

$$\sum_{i=1,2} n_i v_h^i h_q^i + \lambda \sum_{i=1,2} n_i \left(z_i + t \frac{\partial z^i}{\partial q} \right) + \gamma \left(v_h^2 h_q^2 - \widehat{v}_h^2 \widehat{h}_q^2 \right) = 0 \tag{3}$$

where the 'hat' refers to a type-2 household who is mimicking a type-1. Multiplying (1) by z_1, (2) by z_2, and adding both equations to (3), we immediately obtain the A-S Theorem:

$$t \sum_{i=1,2} n_i \frac{\partial \widetilde{z}^i (q)}{\partial q} = \frac{\gamma}{\lambda} \widehat{v}_h^2 \widehat{h}_c^2 \left(z_1 - \widehat{z}_2 \right) = 0 \tag{4}$$

where we have used the envelope condition on q from the second-stage of the household's problem, $h_q^i + g_x^i z_i = 0$, which also applies to the mimicking type-2 household. The function $\widetilde{z}^i (q)$ represents the compensated demand for z_i, where the compensation takes the form of disposable income. The second equality follows from the fact that type-1 households and the mimicking type-2 households have the same disposable income c_1, but differ in their labor supplies. By separability, they will consume the same bundle of goods, so $\widehat{z}_2 = z_1$. Therefore, when the income tax is being set optimally, $t = 0$, so no differential commodities taxes should be applied. This demonstrates the A-S Theorem.

Equation (4) also implies, using the negativity of the own substitution effect, that $t > 0$ if $z_1 > \widehat{z}_2$. This occurs if z and leisure are complements: type-2 mimickers take more leisure than type-1's given that their incomes are the same (Edwards et al, 1994; Nava et al, 1996). It is useful to give the intuition for this result. Starting at $t = 0$, an incremental increase in t will reduce the welfare of a type-i household by z_i. If the income tax schedule is adjusted so that $\Delta c_i = -z_i \Delta t$, there will be no change in utilities v^1 and v^2 and the government budget will remain balanced. But, the mimicker will be worse off implying that the self-selection constraint becomes slack and social welfare can be improved. The same logic applies to the results obtained in the following two sections.

Next we turn to two sorts of extensions to the above analysis. In the first, taken up in the following section, we modify the manner in which goods enter the subutility function $g(x,z)$ by allowing households to have different basic needs or, equivalently, different endowments of one of the goods. In the subsequent section, we consider different specification for labor supply. In each case, the results depend on the relationship between the demands for z by the mimicker and the household being mimicker ($\widehat{z}_2$ and z_1 in the above case).

3. Needs and Endowments

Suppose that, in the manner of the Stone-Geary utility function, households have some basic non-discretionary expenditures that must be spent on one of the goods, say, z. The separable utility function can then be written $u(g(x, z-b), \ell)$, where b is non-discretionary spending on z. One interpretation that can be given to b, following Rowe and Woolley (1999), is that of a basic *need* for good z, such as sustenance, health spending, etc. Alternatively, b might be interpreted as an initial endowment, as in Cremer et al (2001), in which case it takes a negative value.[7] The only difference between the two approaches is that initial endowments enter the overall resource constraint of the economy by adding to net output.[8] Note that b might enter into the utility function in other ways, such as multiplicatively, and a need parameter might be associated with good x as well. Since these would not affect our basic results, we analyze only the case of additive non-discretionary expenditures in good z for simplicity.

If b were the same for all persons, it would obviously have no effect on the A-S Theorem derived in the previous section. The non-discretionary spending would simply be an element of the common utility function faced by all households, which would remain separable. Instead, we assume that b can differ across households. For expositional purposes, we assume that b can take on two values b_j, $j = 1, 2$.

This implies that there can now be four household types, $\{w_i, b_j\}, \ \ i, j = 1, 2$. In analyzing government policy, two informational settings are considered. In one, following Cremer et al (2001), the government can observe neither w nor b. In the other, the government can observe b, but not w. This is the assumption adopted by Rowe and Woolley (1999) in their analysis of needs. We consider these two settings in turn, focusing on the case where b reflects needs rather than endowments.

3.1 Government Does Not Observe Needs

With both w and b unobservable, the government faces a two-dimensional screening problem. This is the case analyzed by Cremer et al (2001). As is well known, the analysis is complex and the results ambiguous, mainly because the directions in which the various self-selection constraints bind are ambiguous. We can simplify the analysis considerably without affecting the main results by assuming that each ability-type is associated with a given need. Thus, a household with wage w_i has a need of b_i. This leaves us with at most one binding self-selection constraint which, unlike in the previous section, can bind in either direction even under a utilitarian objective function. For example, if high-wage households also have high needs, the government may want to redistribute from the low-wage to the high-wage types. For our purposes, that does not affect the issue at hand — whether or not the A-S Theorem applies — although it does affect the consequences if it does not apply. We proceed by assuming that the self-selection constraint applies downwards as in the previous section.

As before, we adopt a two-stage procedure, assuming that in the first stage, labor supply and income are chosen, while in the second stage, disposable income is allocated between the two goods. The analysis of the second stage is identical to earlier. A type-i household chooses z_i to maximize $g\left(c_i - qz_i, z_i - b_i\right)$. This yields the demand function $z(q, c_i, b_i)$, and the value function $h(q, c_i, b_i)$. The envelope theorem for q again yields $h_q^i = -g_x^i z_i$.

In the first stage, the Lagrangean expression for the government's choice of $\{c_i, y_i, q\}$ is exactly as before, and the first-order conditions on c_i and q can be used along with the envelope condition on g to obtain the analog of (4):

$$t \sum_{i=1,2} n_i \frac{\partial \widetilde{z}^i(q)}{\partial q} = \frac{\gamma}{\lambda} \widehat{v}_h^2 \widehat{h}_c^2 (z_1 - \widehat{z}_2) \tag{5}$$

Unlike in the previous section, the right-hand side is generally not zero: it will only be so if $b_1 = b_2$. It can be shown that $t > 0$ if $b_2 > b_1$, and vice versa. That is, if high-wage households also have higher needs, the tax on z should be higher (assuming, of course, that the self-selection constraint on the high-wage types is binding).

This result can be illustrated using Figure 1, which depicts preferences and the budget constraint for a type-1 person and a mimicking type-2. Define the net (after-needs) consumption of good z by $\overline{z}_i \equiv z_i - b_i$. Then, the sub-utility function for the two types of individuals is identical in x and $\overline{z}$, and preferences over x and $\overline{z}$ are independent of labor supply. When $t = 0$, the budget constraints for each household are given by $c_i - b_i = x_i + \overline{z}_i$. The figure shows the choices of x_i and $\overline{z}_i$ for the two types of households when $b_2 > b_1$. As can be seen, $\overline{z}_1 - \widehat{\overline{z}}_2 < b_2 - b_1$ (recall that the 'hat' refers to the mimicker), which implies that $\widehat{z}_2 > z_1$. Since the mimicker purchases more of good z, the self-selection constraint can be weakened by imposing a tax on that good by the same intuition as before. The opposite result occurs if $b_1 > b_2$. Then, $z_1 > \widehat{z}_2$ and it is optimal to set $t < 0$.

As mentioned, if the higher productivity workers have higher needs, the self-selection constraint could apply in the other direction. This is surely the case when productivity differences are very small relative to differences in needs. Then, the mimicker is of type 1, and $z_2 > \widehat{z}_1$. It is optimal to *subsidize* z in this case to relax the self-selection constraint. With four types, the pattern of self-selection constraints becomes quite complex, but as shown by Cremer et al. (2001) the case for a non-zero tax, positive or negative, is very strong.

The upshot of this discussion is that if persons have different needs or endowments, the A-S Theorem will fail to be satisfied even if the utility function is weakly separable.

3.2 Government Observes Needs

Suppose now that the government can observe household needs b_j, but it cannot observe wage rates w_i. There are now four household types, and we denote the government's policy instruments by $\{c_{ij}, y_{ij}, q\}$, $i, j = 1, 2$. However, since needs are now observable, the population can be divided into the two identifiable need types $\{w_1, b_1; w_2, b_1\}$ and $\{w_1, b_2; w_2, b_2\}$. The second stage of the household problem is analogous to above, the only difference being that household demands and functions are now indexed by ij rather than simply i.

The government can now condition its policies on need, and that simplifies matters considerably. In particular, it need worry only about incentive compatibility within each need type. The Lagrangean expression for the government can be written:

$$\mathcal{L} = \sum_i \sum_j n_{ij} v^{ij} \left(h\left(q, c_{ij}\right), \frac{y_{ij}}{w_i} \right) + \lambda \sum_i \sum_j n_{ij} \left(y_{ij} + t z^{ij}\left(q, c_{ij}\right) - c_{ij} \right)$$
$$+ \sum_j \gamma_j \left[v^{2j} \left(h\left(q, c_{2j}\right), \frac{y_{2j}}{w_2} \right) - v^{2j} \left(h\left(q, c_{1j}\right), \frac{y_{1j}}{w_2} \right) \right]$$

The first-order conditions on disposable income and q are:

$$n_{1j} v_h^{1j} h_c^{1j} - \lambda n_{1j} \left(1 - t \frac{\partial z^{1j}}{\partial c_{1j}} \right) - \gamma_j \widehat{v}_h^{2j} \widehat{h}_c^{2j} = 0 \qquad j = 1, 2$$

$$n_{2j}v_h^{2j}h_c^{2j} - \lambda n_{2j}\left(1 - t\frac{\partial z^{2j}}{\partial c_{2j}}\right) + \gamma_j v_h^{2j}h_c^{2j} = 0 \qquad j = 1,2$$

$$\sum_i \sum_j n_{ij}v_h^{ij}h_q^{ij} + \lambda \sum_i \sum_j n_{ij}\left(z_{ij} + t\frac{\partial z^{ij}}{\partial q}\right) + \sum_j \gamma_j \left(v_h^{2j}h_q^{2j} - \widehat{v}_h^{2j}\widehat{h}_q^{2j}\right) = 0$$

It should be apparent that by combining these conditions, we obtain the analog of (4) derived earlier:

$$t\sum_i \sum_j n_{ij}\frac{\partial \widetilde{z}^{ij}(q)}{\partial q} = \sum_j \frac{\gamma_j}{\lambda}\widehat{v}_h^{2j}\widehat{h}_c^{2j}\left(z_{1j} - \widehat{z}_{2j}\right) = 0$$

The last equality comes about because within each need group j, a type-1 household and a type-2 mimicker has the same disposable income and the same value of b_j, so by separability, they have the same demand for good z, or $z_{1j} = \widehat{z}_{2j}$. Therefore, the A-S Theorem applies in this case. It ought also to be obvious that this result extends to other formulations of need, such as multiplicative. Provided the government can classify households by need, and utility functions are separable, the A-S Theorem applies.

The optimal income tax system is obviously more complicated in this case, since there is a different schedule for each need type.

4. Multiple Forms of Labor

In this section, we consider the robustness of the A-S Theorem when household labor supplies are disaggregated into more than one type. For simplicity, we assume that each household can supply two types of labor, say ℓ_c and ℓ_d, whose interpretations will be discussed for various cases considered below.[9] As in the previous section, the informational restrictions that face the government will be key in determining whether the A-S Theorem applies.

4.1 Two Types of Market Labor

Suppose each household supplies two types of labor to the market, receives a wage rate for each, and uses the proceeds to purchase goods. In this case, the utility function becomes $u\left(g\left(x,z\right),\ell_c,\ell_d\right)$. The two types of labor supply could be two different jobs, or the problem could be given an intertemporal interpretation, with ℓ_c and ℓ_d interpreted as present and future labor supply (where x and z can then be interpreted as present and future consumption). In this case, the applicability of the A-S Theorem depends on whether or not incomes from the two forms of labor supply y_c and y_d are observable.

If both y_c and y_d are observable either individually or in the aggregate, the analysis of Section 2 goes through with virtually no modification. The government's selection of an optimal tax policy involves selecting consumption levels and disposable incomes for the two types of households, as well as the commodity tax on z. The conditions on c_i and q are the same as before. Moreover, since the mimicker has the same disposable income as a type-1 person, separability ensures that $\widehat{z}_2 = z_1$, so the optimal commodity tax rate is zero ($t = 0$).

On the other hand, suppose that, say, y_c is observable, but y_d is not. This might correspond with the case in which labor supply ℓ_d is to the underground economy, as in Boadway et al (1994). Of course, for this interpretation to apply, one ought to model explicitly the penalty and detection technologies associated with the underground sector. However, that would serve only to complicate the story without affecting the main result. That result is that the A-S Theorem generally no longer applies if one source of income is not observable to the government.

The intuition is straightforward, even without a formal analysis. If only y_c is observable, the government can control only that part of disposable income that comes from ℓ_c. Assuming that the wage rate of the mimicker is higher that that of a type-1 household in the unobserved sector, it will generally be the case that $\widehat{y}_{d2} \neq y_{d1}$. That implies that $\widehat{c}_2 \neq c_1$, so that even with separable preferences, $\widehat{z}_2 \neq z_1$. So, for example, if $\widehat{y}_{d2} > y_{d1}$ because of the higher productivity of a type-2 person, $\widehat{z}_2 > z_1$, and it will be optimal to impose a tax on good z to relax the self-selection constraint.

More generally, suppose the subutility function $g(x,z)$ is homothetic. In this case, the proportions in which the two goods are consumed by the two persons will be the same. Even in this case it will be optimal to tax good z. In fact, as Boadway et al (1994) show, the optimal commodity tax system is a proportional one on the two goods x and z. The point is that in the absence of full observability of income, a proportional income tax is no longer a perfect substitute for a proportional commodity tax. In the optimum, there needs to be a mix of the two taxes.

4.2 Household Production

Suppose that the second form of labor supply ℓ_d represents non-market or household production with no disposable income that can be used to purchase x and z. All disposable income comes from y_c, which is observed by the government. Assume that the utility function still takes the form $u(g(x,z), \ell_c, \ell_d)$, where the argument ℓ_d incorporates both the disutility of the non-market work as well as the the product of that work. In this case, the A-S Theorem still holds regardless of whether non-market labor is observed by the government. Indeed if all households share the same preferences, that will not be relevant.

The analysis is a straightforward application of that used in Section 2. The government controls y_c, and therefore disposable income that is used to purchase x and z. The existence of non-market labor complicates things slightly because it conditions the structure of the optimal non-linear income tax, and might in principle affect the direction in which the incentive constraint is binding. Suppose, for example, that the self-selection constraint is binding on type-2 households in the optimum. Households of type 2 who mimic those of type 1 will earn the same market income y_{c1} and obtain the same disposable income c_1. By the same analysis as above, $\widehat{z}_2 = z_1$, and so $t = 0$ in the optimum. This logic still applies if the self-selection constraint binds on type-1 households.

4.3 Different Preferences for Leisure

Household might differ not only by ability but according to their preferences for leisure. This adds another important and difficult dimension to redistributive policy. For one thing, governments are unlikely to be able to differentiate persons according to their preferences for leisure, that is, their laziness or diligence. For another, even if they could, it is not obvious how redistributive policies ought to differentiate among preference types. There is a school of thought that suggests that households are responsible for their own preferences, and redistributive policies ought only to compensate for ability differences (Roemer, 1995; Fleurbaey and Maniquet, 1999). On the other hand, as stressed by Cuff (2000), preferences for leisure might be viewed as being partly determined not just by one's attitude to work, but also to the degree of difficulty individuals face in working.

Consider the simple case in which there are two ability-types of households and two preference types. A convenient way to formulate the utility function when there are differences in preferences is as $u\left(g\left(x,z\right),\alpha\ell\right)$.[10] In this formulation, α can take on the values α_1 and α_2. If $\alpha_1 > \alpha_2$, preference type-1 has a greater preference for leisure than preference type-2 households. In the unlikely event that the government could distinguish between households with high and low preference for leisure, it could simply design two separate non-linear income tax systems for the two types, exactly analogously to the case of different needs for goods considered earlier. In this case, it is obvious that the A-S Theorem applies, since within each preference type, the high-wage mimicking person would have the same disposable income as the low-wage person, and by separability would consume the same bundle of goods.

If the government cannot distinguish preference types, it is again faced with a two-dimensional screening problem. Depending on the relative welfare weights attached to the two preference types, the pattern of binding self-selection constraints can vary (Boadway et al, 2001). However, regardless of what the pattern might be in the optimum, the separability of the utility function combined with the commonality of the sub-utility function $g(x, z)$ implies that the A-S Theorem still holds. Mimickers will have the same income and disposable income as those they are mimicking regardless of the type of either. Therefore, they will consume the same bundle of goods, implying that a differential commodity tax cannot be used to separate the two types. Differences in preferences for leisure merely serve to complicate the form of the optimal non-linear income tax.

Finally, note that differences in preference for leisure could reflect differences in need, analogous to the case of differences of need for different goods. For example, utility functions might take the form $u(g(x, z), \ell + a)$, where a reflects need and can vary from one household to another. By similar reasoning to above, the A-S Theorem continues to apply in this case regardless of whether the government can observe household needs.

4.4 Becker-Gronau Household Production

A final case to consider is the case where consumption of goods itself requires the allocation of some time, following Becker (1965), Gronau (1977) and Jacobsen Kleven (2000). One way of formulating the utility function in this case is as

$u(g(X,Z),\ell)$, where X and Z are commodities produced by household production functions $f^x(x,\ell_x)$ and $f^z(z,\ell_z)$, where ℓ_x and ℓ_z are labor inputs into the production of the home-produced commodities. Assuming that the household production functions are the same for both households, the A-S Theorem applies directly. Mimicking households will have the same income and disposable incomes of those being mimicked, and given these assumptions will purchase the same quantities of the two goods x and z to produce the same quantities of household goods. On the other hand, if the households had different productivities in home production, which might be a reasonable assumption, the same disposable incomes would generally give rise to different demands for x and z by mimicking type-2's and type-1's.

5. The Linear Progressive Income Tax Case

The A-S Theorem represented a landmark in optimal tax theory. It showed, in the context of a standard optimal income tax setting, that there is no need to supplement the income tax with differential commodity taxes when goods are weakly separable from leisure in household preferences. This result was derived under what might now be regarded as the strong assumption that all households have the same preferences. We have relaxed this assumption in two realistic ways: first, that needs for particular goods can differ across households, and second, that households may supply more than one form of labor. In each of these cases, the A-S Theorem will generally be violated if the government is unable to observe the difference in needs or some types of labor supply.

These results, like the A-S Theorem itself, presume that the government has full freedom to levy an optimal non-linear income tax. It is straightforward to generalize the analysis to a more restrictive income tax regime, such as an optimal linear income tax. In this case, it is well known that differential commodity taxes should not be used if goods are separable from leisure and if the sub-utility function in goods is quasi-linear, that is, Engel curves for x and z are linear (Deaton, 1979).

By using the same logic as before, differences in need and endowments for goods will have the same effect on the applicability this modified A-S Theorem, as in the A-S Theorem under non-linear optimal income taxation. That is, if households have different unobservable needs for one of the goods, it will generally be desirable to impose a tax or subsidy on it. On the other hand, if needs are observable, different tax schedules will apply to persons of different needs classes. Similarly, if households supply two types of labor, the modified A-S Theorem applies if both types are observable, but not otherwise. As well, the modified version still applies if there is unobserved household production or if preferences for leisure differ.

6. Conclusions

When looking at actual tax systems, one finds almost everywhere a mix of direct and indirect taxes, or more precisely of consumption and income taxes. What are the reasons for such an apparent violation of the A-S Theorem? Ignorance of basic public economics and thus bad fiscal engineering? Huge compliance costs in income taxation relative to consumption taxation? Reasons developed in this paper and elsewhere for infirming the A-S Theorem? Unwillingness to implement an optimal income tax? Lack of separability of the utility function?

As usual the answer is ‘a bit of everything’. It is clear that in developing countries, the compliance and administration costs of income taxation are so high that tax authorities have to rely on friendlier indirect taxation. The arguments developed above have also some empirical relevancy. For example, it is natural to think that needs differ across individuals and are not always observable. The issue of separability is also far from being settled: most econometric studies do not lend support to such separability. It is possible that some public finance experts and policy makers miss the point of A-S Theorem and believe that taxes are like eggs: you do not put them in the same basket as income taxation. Finally, there is an issue with the willingness to implement an optimal income tax. The A-S Theorem assumes that one starts with such a tax. It is far from being granted that existing income tax systems correspond to such a scheme, and without optimal income taxation there is no A-S Theorem.

ENDNOTES

Acknowledgements: This paper, like much of our work in public economics, is inspired by Joe Stiglitz. His combination of analytical virtuosity, intuition, and dedication to economic science as a vehicle for addressing pressing social concerns is an example to us all. We are grateful to Richard Arnott for comments on an earlier draft, and to our various collaborators — especially Helmuth Cremer, Michael Keen, Maurice Marchand and Nicolas Marceau — on whose work and ideas we have drawn.

1. See Cremer et al. (2001) and references cited therein. See also Naito (1999), who shows that if production consists of several sectors using in variable proportion the different types of workers, then it pays to tax the sectors employing a relatively high proportion of skilled labor; Saez (2002), who shows that Naito's objection disappears in the long run; and Cremer and Gahvari (1995), who underline the desirable insurance effect of commodity taxation.

2. This assumption assures that the A-S Theorem applies. It has been questioned on empirical grounds. See Browning and Meghir (1991).

3. Revesz (1986) has shown that if the government could levy license fees alongside proportional commodity taxes, it might be optimal to do so even if optimal proportional commodity tax rates are zero because of separability. In this

paper, we assume that license fees cannot be enforced because of the possibility of resale.

4. A similar procedure has been used by Edwards et al (1994), Nava et al (1996), and Cremer et al (2001).

5. In what follows, variables applying to households of type i are denoted by a subscript, while functions for household i are denoted by a superscript. Function subscripts refer to partial derivatives.

6. The first-order conditions on incomes y_i can be used to characterize the structure of the optimal income tax. The characterization is standard, and we suppress it here. In what follows, the government is always taken to be applying the optimal income tax.

7. An interpretation of unobserved endowments that gives rise to a rationale for differential commodity taxation is the case of bequests, analysed in Boadway et al (2000) and Cremer et al (2002). In this case, the analysis is intertemporal, and the differential taxation applies to future versus present consumption. Capital income taxation then becomes the policy instrument for taxing future consumption differentially.

8. Differences in needs or endowments are also similar to heterogeneous tastes. On this, see Saez (2000).

9. This is different from the case where different households supply different types of labor. If these are not substitutable, the A-S Theorem may also fail as shown by Naito (1999), noterd earlier.

10. This is the formulation for the preferences for leisure used by Boadway et al (2002), who study the design of the optimal redistributive income tax when households differ in both ability and preferences.

REFERENCES

Atkinson, Anthony B. and Joseph E. Stiglitz. "The Design of Tax Structure: Direct versus Indirect Taxation". *Journal of Public Economics* 6 (1976): 55–75.

Becker, Gary S. "A Theory of the Allocation of Time". *The Economic Journal* 75 (1965): 493–517.

Boadway, Robin, Maurice Marchand and Pierre Pestieau. "Towards a Theory of the Direct-Indirect Tax Mix". *Journal of Public Economics* 55 (1994): 71–88.

Boadway, Robin, Maurice Marchand and Pierre Pestieau. "Redistribution with Unobservable Bequests: A Case for Taxing Capital Income". *Scandinavian Journal of Economics* 102 (2000): 253–67

Boadway, Robin, Maurice Marchand, Pierre Pestieau and Maria del Mar Racionero. "Optimal Redistribution with Heterogeneous Preferences for Leisure". *Journal of Public Economic Theory* (2002), forthcoming.

Browning, Martin and Costas Meghir. "The Effects of Male and Female Labor Supply on Commodity Demands". *Econometrica* 59 (1991): 925–951.

Corlett, W.J. and D.C. Hague. "Complementarity and the Excess Burden of Taxation". *Review of Economic Studies* 21 (1953–54): 21–30.

Cremer, Helmuth and Firouz Gahvari. "Uncertainty and Optimal Taxation: In Defense of Commodity Taxes". *Journal of Public Economics* 56 (1995): 291–310.

Cremer, Helmuth, Pierre Pestieau, and Jean-Charles Rochet. "Direct Versus Indirect Taxation: The Design of the Tax Structure Revisited". *International Economic Review* 42 (2001): 781–99.

Cremer, Helmuth, Pierre Pestieau and Jean-Charles Rochet. "Capital Income Taxation when Inherited Wealth is not Observable". *Journal of Public Economics*

(2002), forthcoming.

Cuff, Katherine. "Optimality of Workfare with Heterogeneous Preferences". *Canadial Journal of Economics* 33 (2000): 149–74.

Deaton, Angus. "Optimally Uniform Commodity Taxes". *Economics Letters* 2 (1979): 357–361.

Edwards, Jeremy, Michael Keen and Matti Tuomala. "Income Tax, Commodity Taxes and Public Good Provision: A Brief Guide". *Finanzarchiv* 51 (1994): 472–97.

Fleurbaey, Marc and Francois Maniquet. "Compensation and Responsibility". (1999), mimeo.

Gronau, Ruben. "Leisure, Home Production and Work — The Theory of the Allocation of Time Revisited". *Journal of Political Economy* 85 (1977): 1099–123.

Guesnerie, Roger. *A Contribution to the Pure Theory of Taxation.* Cambridge, UK: Cambridge University Press, 1995.

Jacobsen Kleven, Henrik. "Optimal Taxation and the Allocation of Time". (2000), mimeo.

Mirrlees, James A. "An Exploration in the Theory of Optimal Income Taxation". *Review of Economic Studies* 38 (1971): 175–208.

Naito, Hisahiro. "Re-examination of Uniform Commodity Taxes under a Nonlinear Income Tax System and its Implication for Productive Efficiency". *Journal of Public Economics* 71 (1999): 165–188.

Nava, Mario, Fred Schroyen and Maurice Marchand. "Optimal Fiscal and Public Expenditure Policy in a Two-Class Economy". *Journal of Public Economics* 61 (1996): 119–37.

Revesz, John T. "On Some Advantages of Progressive Indirect Taxation." *Public*

Finance XXXXI (1986): 182–99.

Roemer, John E. *Equality of Opportunity.* Cambridge, US: Harvard University Press, 1998.

Rowe, Nicholas and Frances Woolley. "The Efficiency Case for Universality". *Canadian Journal of Economics* 32 (1999): 613–29.

Sandmo, Agnar. "Optimal Taxation: An Introduction to the Literature". *Journal of Public Economics* 6 (1976): 37–54.

Saez, Emmanuel. "The Desirability of Commodity Taxation under Nonlinear Income Taxation and Heterogeneous Tastes". NBER Working Paper (2000), *Journal of Public Economics*, forthcoming .

Saez, Emmanuel. "Direct or Indirect Tax Instruments for Redistribution: Short-run versus Long-run". NBER Working Paper, (2002), *Journal of Public Economics*, forthcoming.

Stiglitz, Joseph E. "Self-Selection and Pareto Efficient Taxation". *Journal of Public Economics* 17 (1982): 213–40.

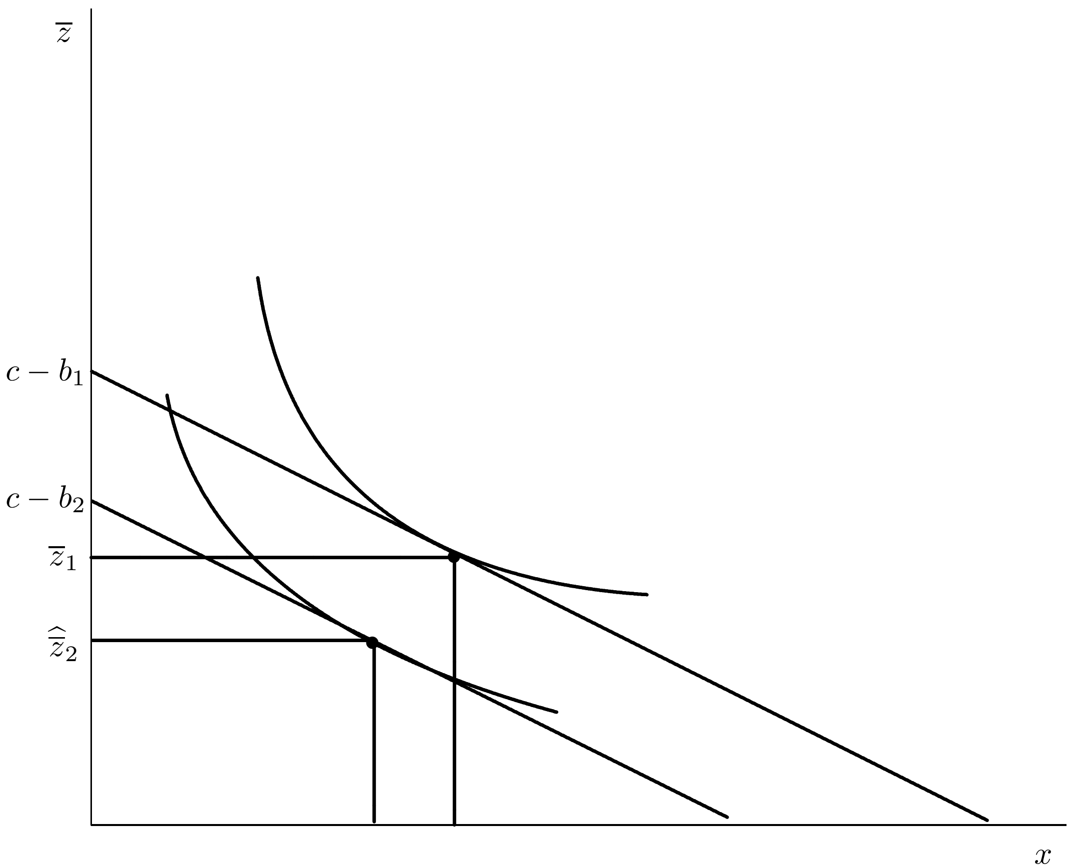

Figure 1