The limitations of decentralized world redistribution: An optimal taxation approach[☆]

Wojciech Kopczuk[a], Joel Slemrod[b,*], Shlomo Yitzhaki[c]

[a] *Columbia University, New York, NY 10027, USA*
[b] *Department of Economics, University of Michigan, 611 Tappan Street,*
Ann Arbor, MI 48109-1220, USA
[c] *Hebrew University of Jerusalem, Jerusalem, Israel*

Abstract

A centralized scheme of world redistribution that maximizes a border-neutral social welfare function, subject to the disincentive effects it would create, generates a drastic reduction in world consumption inequality, dropping the Gini coefficient from 0.69 to 0.25. In contrast, an optimal decentralized (i.e., with no cross-country transfers) redistribution has a miniscule effect on world income inequality. Thus, the traditional public finance concern about the excess burden of redistribution cannot explain why there is so little world redistribution.

Actual foreign aid is vastly lower than the transfers under the simulated world income tax, suggesting that voluntary world transfers – subject to a free-rider problem – produces an outcome that is consistent with rich countries such as the United States either placing a much lower value on the welfare of foreigners, or else expecting that a very significant fraction of cross-border transfers is wasted. The product of the welfare weight and one minus the share of transfers that are wasted constitutes the implicit weight that the United States assigns to foreigners. We calculate that value to be as low as 1/2000 of the value put on the welfare of an American, suggesting that U.S. policy is consistent with social preferences that place essentially no value on the welfare of the citizens of the poorest countries, or that implicitly assumes that essentially all transfers are wasted.
© 2003 Elsevier B.V. All rights reserved.

---

[*] Corresponding author. Tel.: +1-734-936-3914; fax: +1-734-763-4032.
  E-mail address: jslemrod@umich.edu (J. Slemrod).

1. Introduction

In May, 2002 rock star Bono and the former U.S. Secretary of the Treasury Paul O'Neill toured Africa together. At each stop they publicly aired their different views on the need for and effectiveness of foreign aid. Bono insisted that more aid is needed to lift Africa out of desperate poverty, implying that it is largely the mendacity of developed countries that prevents more aid. Secretary O'Neill argued that much aid has done little to reduce poverty, owing in large part to waste and corruption.

This high-profile tour generated wide media coverage of global poverty and global income inequality. But the same debate has been ongoing for many years. Gross disparities of income across countries[1] have drawn attention to the small amount of resources transferred from the rich countries of the world to the poor countries, and have given rise to calls that the rich countries devote much more of their resources to foreign aid. For example, Sachs (2001) has called for the United States to double its aid budget and devote the funds to disease control, primary education, clean water, and other vital needs of impoverished places.

That voluntary, decentralized cross-country transfers might fall far short of what a centralized scheme motivated by the pursuit of social justice would imply is no surprise. To the extent that the well-being of the poor is a public good from which residents of all rich countries derive pleasure, the equilibrium outcome will be inefficient because of the free-rider problem. Alternatively, some donors might derive utility not only from the well-being of the recipient, but also from the act of giving itself.

The unwillingness of the United States and other developed countries to substantially raise their foreign aid may also reflect one or both of two other factors: The citizens of rich countries place a very low value on the welfare of the citizens of poor countries, or they may shy away from transfers because of the large efficiency cost that would plague such efforts. This cost may have two sources. One is the concern expressed by Secretary O'Neill and others that the funds would not reach the targeted groups due to waste and corruption. Another type of cost relates to the traditional concern of public finance economists that the process of taxing the well off and transferring the proceeds to the less well off causes disincentives. The economic cost of these disincentives limits the optimal amount of cross-country transfers that would be undertaken even by a policymaker with egalitarian impulses to redistribute from the globally rich to the globally poor.

In this paper we adapt the predominant perspective in public finance by presuming that the rationale for redistribution is neither altruism nor the "warm glow" of giving, but rather the pursuit of social justice. Following this tradition, we assume that the social desirability of alternative distributions of welfare among the world's population can be assessed in terms of a Bergson–Samuelson "social welfare function" defined over the utility (further assumed to be a function only of income) of the population under study. This function is assumed to be both anonymous and Paretian. From this public finance perspective, it is clear that the problem of global redistribution has

---

[1] Milanovic (2002) has shown that the major source of inequality among the people of the world is cross-country differences.

the same structure as the problem each country faces – trading off in the pursuit of social justice the efficiency costs of a progressive tax system against the more equal distribution of welfare it achieves. In fact, most countries achieve some degree of redistribution through their own tax-and-transfer system. Clearly, the extent of overall world redistribution is small relative to world inequality because cross-country transfers are minimal. The question of whether these minimal transfers are at least approximately *optimal* and what the optimal transfers would be requires further investigation, however.

In this paper we explore this question quantitatively as follows. We first calculate each country's optimal redistributive policy, assuming that each country sets its tax system to maximize a concave social welfare function of individual utility levels, knowing that the tax system will influence individuals' choices. Then each country will set its own tax schedule that is more or less progressive based on the distribution of incomes (more precisely, the ability to earn income) *within that country.* Even though the social welfare function is concave, the desire to redistribute is constrained by the economic cost of the marginal tax rates the redistribution requires. Using data on income inequality and assumptions about utility functions that imply how responsive behavior is to taxation, we calculate the optimal income tax system in each of 118 countries and characterize the amount of redistribution that these decentralized systems produce.

Then we consider the hypothetical case of a world income tax, where the same tax schedule applies to everyone regardless of where they live, and which therefore allows for transfers across countries.[2] We first consider the case where there is no waste (other than excess burden) from cross-country transfers and that the tax setter is *border-neutral*, meaning that each person's welfare enters the social welfare function the same regardless of where he or she lives. Assuming further that the world decision maker has the same preferences as each country about the tradeoff between the mean and distribution of incomes (i.e., an equally concave social welfare function), and faces the same costs from imposing redistribution, we can solve for the optimal progressivity of the world income tax. The solution depends on the inequality of world incomes, and not on the degree of inequality within countries.

The results of simulating these stylized models reveal that the decentralized tax-and-transfer scheme makes hardly any dent in the world income inequality. This is so even though countries pick progressive tax systems on their own. In contrast, an optimal world income tax would significantly reduce the world inequality of consumption, albeit with a larger efficiency cost and at the cost of a reduction in welfare of citizens of the richest 25 countries. Thus, we conclude that a concern about the excess burden of cross-country transfers cannot justify why foreign aid is so low – what limits these transfers is not the efficiency cost of the redistribution.

Using our methodology the limitations of voluntary, decentralized redistribution can be quantified. In the final section of the paper we address this question by allowing

---

[2] One could also consider a version of the centralized tax that allowed each country to have its own marginal tax rate and demogrant. We have not pursued this because of the large computational burden of solving this problem with $2 * n$ instruments, where n is the number of countries. We do know, though, that the solution of the $2 * n$ instruments problem includes as a special case the world income tax solution we calculate (by having all demogrants equal and all tax rates equal). Therefore, we can definitely say that the $2 * n$ solution will result in a reduction in world inequality that is at least as large as we find.

the policy makers in the rich countries to place a lower value on the welfare of the citizens of other countries at any given level of income compared to their own citizens – as suggested by Bono – and/or expect that a fraction of cross-country transfers would be wasted – as suggested by Secretary O'Neill. With our parameter assumptions we cannot distinguish *between* the Bono and O'Neill scenarios, but we can calculate what the product of that relative value and the share of transfers that is wasted must be in order to generate the current level of cross-country transfers, in the form of foreign aid, given by rich to poor countries.

It is shockingly low. In our baseline case, the level of U.S. foreign aid is consistent with foreigners are on average being valued by the U.S. at just 16% of an average American, with the citizens of the poorest countries weighted by as little as 1/20th of 1%. The latter value implies either that voluntary, decentralized aid produces an outcome that is consistent with the U.S. putting essentially no weight on the welfare of those individuals, or that 1999/2000th of the transfer is wasted, or a combination of both. One reason for this low implicit valuation is that donor countries may view the welfare of the poorest foreigners as a public good, and so voluntary donations are subject to a free-rider problem.

2. Methodology

2.1. Calculating the optimal linear income tax

Our central analytical tool is a model of the optimal income tax structure, as pioneered by Mirrlees (1971). The idea is that the government chooses an income tax function that maximizes a given social welfare function, subject to an exogenously specified revenue requirement and the constraint that individuals will choose the levels of consumption and leisure that maximize their utility subject to their own budget constraints, which depend on the tax system chosen.

There are three key elements of the problem. The first is the degree of concavity of the social welfare function, which captures how society makes the tradeoff between the sum of utilities and the distribution of utilities. Second is the elasticity of substitution between leisure and consumption in individuals' utility functions (which are assumed to be identical); this determines the amount of distortion, or welfare cost, for any given tax structure. The final element is the distribution of abilities, where an individual's ability is presumed to be equal to the pre-tax wage rate. Loosely speaking, the optimal income tax structure trades off the social welfare gains of a more equal distribution of utilities against the efficiency cost caused by the structure of marginal tax rates needed to achieve any given amount of redistribution. [3]

---

[3] The standard optimal progressivity problem that we draw on is atemporal, and therefore does not directly address the effect of large transfers on the incentives to save and invest; this could be an additional cost of redistribution. This is especially noteworthy given that growth models were an important theoretical underpinning of modern foreign aid, suggesting that transfers could overcome failures in international capital markets and increase investment and, therefore, growth. That aid might serve to increase consumption rather than investment might be viewed as wasteful by donor nations, a subject we elaborate on in Section 4.

Although the optimal income tax literature has explored the sensitivity of the results to various assumptions about the social welfare function, the distribution of abilities, and the magnitude of behavioral response, it has not been used to quantitatively explore the implications of a decentralized system of redistribution in a world of gross inequalities across countries. This is the task we begin below.

2.2. Choosing the model parameters

There are two scenarios that we wish to compare. One is a decentralized solution, in which each country selects its own optimal linear income tax system. The other one is a world income tax system, in which the decision maker designs a single linear income tax that applies to all individuals in the world. This exercise requires making a host of assumptions about the distribution of earning potential, the utility function, welfare function, behavioral elasticities and stylized economies we study. In what follows we review the main issues.

2.2.1. The distribution of abilities

The dispersion of abilities is critical because, in general and ceteris paribus, the optimal linear income tax will be more progressive (i.e., feature a higher demogrant and higher tax rate) the more unequal is the initial distribution of earning potential within the jurisdiction. Mirrlees (1971) presents an example in which widening the distribution of skills, assumed equal to wage rates, increased the optimal marginal tax rates; he concludes that the dispersion of skills necessary to imply marginal tax rates much higher than the 20–40% range is unrealistically high. In his baseline numerical simulation, he sets the value of the standard deviation of the associated normal distribution (denoted σ) in the assumed logarithmic distribution of skills to be equal to 0.39, derived from Lydall's (1968) figures for the distribution of income from employment in various countries. When Mirrlees repeated the simulation with $\sigma = 1.0$, a much wider dispersion of ability, he reported that the optimal tax schedule

> "is in almost all respects very different. Tax rates are very high: A large proportion of the population is allowed to abstain from productive labour. The results seem to say that, in an economy with more intrinsic inequality in economic skill, the income tax is a more important weapon of public control than it is in an economy where the dispersion of innate skills is less. The reason is, presumably, that the labour-discouraging effects of the tax are more important, relative to the redistributive benefits, in the latter case."

Stern (1976), examining only flat-rate tax systems, corroborates Mirrlees finding. For his base case featuring an elasticity of substitution between goods and leisure of 0.4, when $\sigma = 0.39$, the optimal marginal tax rate is 0.225, but it rises to 0.623 when $\sigma = 1$. Cooter and Helpman (1976) perform a variety of numerical simulations, and find that

for all of them the optimal marginal tax rate increased as the constant-mean ability distribution spreads out.[4]

Of course, innate ability is unobservable, so its dispersion is not knowable, either. What is available, and is collected in Deininger and Squire (1996), are estimates of Gini coefficients for 138 countries. These estimates were produced from a variety of micro data sources, and come from studies of varying quality. They identify Gini coefficients based on actual observation of individual units drawn from household surveys, based on comprehensive coverage of the population, and based on comprehensive coverage of different income sources as well as of population groups. World Bank (2000, Table 2.8) is a more recent source of Gini coefficients. These estimates are based on survey data obtained from government statistical agencies and World Bank country departments, and in many cases overlap with the Deininger and Squire (1996) observations. In our simulations, we use the World Bank (2000) estimates as the primary source, and resort to the "high-quality" observations in Deininger and Squire (1996) for countries that are not present in that data set.

A more vexing problem is that the studies sometimes calculate the inequality of pre-tax income, sometimes calculate the inequality of after-tax income, and sometimes calculate the inequality of consumption. Of course, none calculates the inequality of ability. By making strong assumptions about the process that generates income, one could claim to have recovered the distribution of abilities that is consistent with the data. For example, for a given and common utility function and tax system, one could convert the distribution of labor earnings into the distribution of abilities. This is the procedure we follow.

Because of the greater variability of annual income compared to annual consumption, measures of inequality based on the former will tend to be higher. Deininger and Squire report that in their sample the mean difference between the expenditure-based Gini coefficients and those based on gross income is 0.066. They also report that for the 19 pairs of Gini coefficients computed using the Luxembourg Income study data, those based on after-tax income were on average 3 points lower than those based on gross income; this sample includes, however, only one developing country (Mexico). Clearly, the quantitative importance of this effect will depend on the effective progressivity of the tax system in place.

In what follows we assume that the distribution of abilities in each country is lognormal. Then, we parameterize the distribution so that the resulting Gini coefficient of income or consumption for a given country under a certain baseline income tax system[5] is equal to the empirical value (full details of our calibration methodology are discussed in Section 2.4). In this exercise, gross income is assumed to equal

---

[4] Helpman and Sadka (1978) claim that this result is not general, but offer only a trivial counter-example that features a Rawlsian (maximin) social welfare function and a fixed lowest ability level of zero. They argue that there should exist counter-examples with more general social welfare functions, but admit they were unable to identify any such example.

[5] The baseline income tax system features a marginal tax rate of 0.30.

labor income of the individual, and both consumption and net income are assumed to correspond to after-tax income. [6]

2.2.2. The individual utility function

The individual utility function is a critical element of the problem because it determines the substitutability between leisure and consumption, which in turn reveals the marginal efficiency cost of any degree of tax progressivity. In his simulation analyses of the optimal linear income tax, Stern (1976) focuses on a constant-elasticity-of-substitution (henceforth CES) utility function with an elasticity of substitution of 0.4, based on his reading of the labor supply elasticity literature available at that time. Depending on how it is read, the literature since then suggests considering both a lower and a higher number: Lower because the aggregate elasticity of substitution between leisure and consumption may be less than 0.4, [7] higher because labor supply is only one dimension of behavioral response to taxation that involves an efficiency cost, and research on the elasticity of taxable income suggests that an elasticity of 0.6 may be appropriate (Auten and Carroll, 1999; Gruber and Saez, 2002; Slemrod, 1998). Although in this case the relevant behavioral response is summarized by an elasticity of taxable income rather than an elasticity of substitution between consumption and leisure, in order to be comparable with most of the optimal income tax literature we retain the standard modeling. However, we assign higher behavioral responses than have been found for labor supply, in order to represent the whole range of possible responses.

Somewhat surprisingly, the "income elasticity" of optimal progressivity – do richer countries choose more progressive tax systems? – in this class of models has been almost completely ignored. Indeed, the answer is not obvious. A proportional increase in all individuals' abilities changes the set of tax systems that raise the required amount of revenue. Under certain conditions, the admissible tax systems are simply scaled up in the sense that an equi-proportionate change in all abilities, revenue, and the demogrant, holding the marginal tax rate constant, is still admissible (but perhaps no longer is optimal). However, holding taxes and the degree of inequality constant, the commonly used CES utility functions with an elasticity of substitution below unity imply that in countries with high average ability levels there is much less labor supply, relative to countries with low average abilities, than is apparently observed. As a result, the tax base and revenue collected increase less than proportionally, so that it is not possible to sustain a scaled-up tax system.

One approach to these issues is to consider the class of utility functions that yield a "scale" elasticity of zero. [8] As discussed by King et al. (1982), this class has the

---

[6] This procedure thus does not account for the fact that, due to life-cycle consumption smoothing, the inequality of consumption might be less than the inequality of after-tax income, for a given distribution of abilities. Nor does it attempt to sort out the contribution of capital income and land rent in the distribution of income; data limitations would preclude doing much about these issues, and especially the latter one.

[7] For a survey of the labor supply literature see Blundell and MaCurdy (1999).

[8] Write leisure as $L(sw, sG)$, i.e., a function of wage rate and income, where s is a scalar. The necessary property for a zero scale elasticity is that L is homogenous of degree zero in (w, G), $dL/ds = 0$. Note that this property depends on a combination of income and price responses.

form $U(\ln(C) + g(L))$, where C is consumption and L is leisure. The motivation for examining this utility function is to ensure that simulations yield results that are not grossly inconsistent with the empirical observation that labor supply is broadly similar across countries with widely varying average income levels. Note, though, that the *optimal* tax system may not simply scale up, because the optimum also depends on the social welfare function. What the assumption about utility functions guarantees is that, ceteris paribus, the income elasticity of the optimal tax structure depends only on the social welfare function.

In what follows, we present results for the Cobb–Douglas utility function. This is the only CES utility function that is also in the King–Plosser–Rebelo class. This choice implies a compensated elasticity of labor supply of one, which is high in the context of the literature on the elasticity of taxable income, but within the range of available estimates. The relatively high compensated elasticity means that our estimates of the efficiency cost of high marginal tax rates are higher than otherwise, and thus provide a larger constraint than otherwise on the extent of redistribution that is appropriate.

2.2.3. The social welfare function

Although there have been attempts to recover a society's social welfare function (henceforth SWF) from examining actual government policies, or by examining individual risk aversion, for the most part economists have not tried to defend a particular SWF. Instead, they have investigated the implications of alternative specifications of the SWF for the solution to the problem at hand. We adopt that strategy as well.

To be consistent with the earlier literature, we investigate SWFs of the type introduced by Atkinson (1970), that are of the form $W = \Sigma(1 - v)^{-1}U^{1-v}$. The higher the value of v, the larger is the concavity of the SWF, and the larger is the implied willingness of the society to trade off the sum of utilities for a more equal distribution of the utilities. We investigate the implications of three values of v: 0.5, 2.0, and 5.0, but concentrate on the case of $v = 2.0$, which is Stern's (1976) central case, as well. Whatever value we choose, we assume it is the same for all countries and for the designer of the world income tax. In so doing, we skirt the fascinating but difficult question of whether the degree of egalitarianism differs across countries, including whether it differs systematically depending on the mean level of income or on the distribution of abilities.

2.3. Introducing tradables and PPP differences

In practice, there are significant differences in countries' real price levels. Ignoring these differences would have some peculiar implications when we allow for transfers across countries. The centralized budget constraint would simply add up nominal taxes and subsidies of different economies, so that it would amount to assuming that U.S. and Indian consumption can be exchanged one for one. While this may be correct for tradable commodities, it is not correct for the non-tradable ones. There are also implications for the location of production. Ignoring the presence of non-tradable commodities and holding price levels fixed while allowing for large international transfers will invariably lead to poor countries shutting down their production and relying solely on transfers. The prediction of 100% voluntary unemployment across the Third World would be a highly undesirable model feature.

In this section we develop a model that addresses these issues. The model features two sectors in each country that produce tradable and non-tradable commodities, denoted T and N, respectively. We normalize the (world) price of tradable goods to one. Non-tradable commodities are produced and consumed domestically. Because people want to consume both types of goods, some non-tradable goods have to be produced in each country. Equilibrium is reached by the adjustment of relative wages in the two sectors.[9] The details of the model follow.

2.3.1. Individuals

Assume that there is a continuum of individuals characterized by (heterogeneous) skill levels a. We consider the following utility function:

$$u(T, N, L) = [(1 - \alpha)(T^\delta N^{1-\delta})^{-r} + \alpha L^{-r}]^{-1/r}.$$

This utility function is CES between leisure (L) and consumption commodities. The Cobb–Douglas consumption segment implies that the fraction δ of total income is spent on tradables, while the rest is spent on non-tradables. Denoting the price of non-tradables in country i as p_i, consumption of the two types of goods is therefore given by

$$T^D = \delta(G + (1 - t)w(a)(1 - L)), \quad N^D = \frac{1 - \delta}{p_i}(G + (1 - t)w(a)(1 - L)),$$

where $w(a)$ is the wage rate of an individual with the skill level of a and G is a lump-sum demogrant that may be provided to each individual by the government.

2.3.2. Production

We assume that production in both sectors takes place using only labor. However, the relative productivity of workers with different skill levels varies by sector. Each individual works in just one sector. More specifically, we assume that production in the tradable sector takes place using efficiency units of labor, such that

$$T^s = \int_{S_T} a(1 - L(a)) \, dF(a),$$

where the integration takes place over the set of workers who choose to work in the tradable sector, S_T. The productivity of a worker in the non-tradable sector is assumed to be more closely related to the amount of time that is invested in the activity, although it is positively correlated with skill. In particular, we assume that the productivity in the non-tradable sector is a^d, where $0 \leqslant d < 1$, so that

$$N^s = \int_{S_N} a^d(1 - L(a)) \, dF(a),$$

---

[9] An alternative equilibrating mechanism would allow for the substitution of labor for capital. We do not, however, consider this to be a realistic possibility. For example, we are not aware of a conceivable way of substituting capital for the time of a barber. This example captures an important feature of at least some non-tradable commodities: They require the time of an individual. In other words, in these cases highly skilled individuals are not more productive (or at least they are not much more productive) than the low-skilled ones.

where S_N is the set of workers that choose to work in the non-tradable sector. In the extreme case when $d = 0$, each individual is equally productive in the non-tradable sector. In general, more skilled individuals are more productive in the non-tradable sector, but by a smaller (and decreasing) factor than in the tradable sector. There are no country-specific productivity differentials other than differences in the skill levels of individuals.

2.3.3. Equilibrium

We assume that both sectors are competitive. Because the tradable good is the numeraire, the individual who chooses to work in the tradable sector will receive a wage rate equal to a per unit of his time. The individual who chooses to work in the non-tradable sector is paid $p_i a^d$. Thus,

$$w(a) = \begin{cases} a, & a > p_i a^d, \\ p_i a^d, & \text{otherwise.} \end{cases}$$

Because it is assumed that $d < 1$, low-skilled workers will choose to work in the non-tradable sector and high-skilled workers will choose to work in the tradable sector, although the cutoff level of skill will differ from country to country and depend on the tax system in place.

2.3.4. Features of the equilibrium

The price of the non-tradable commodity determines the potential wages of every individual in each of the two sectors, which determines the sector in which individual is working and allows us to solve for individual consumption and labor decisions. Therefore, the price of non-tradables in a given country is sufficient to determine the aggregate demand and supply of non-tradables. In the equilibrium, the price adjusts to make them equal.[10] Total imports of tradables must be equal to the transfer to the country:

$$T^D - T^S = transfer,$$

because transfers can only take the form of tradables.

One feature of equilibrium is that richer economies have a higher price of non-tradable commodities, so that the overall price level in richer economies is higher. This is also a well-known property of actual relative price levels, remarked upon by Balasa (1964) and Samuelson (1964),[11] who suggest explanations that are in the same spirit as this model.

2.4. Calibration methodology and baseline results

Table 6 lists the key data for all of the 118 countries we examine. The first column lists the population in 1999. Note that, although not all countries in the world

---

[10] The level of inequality may affect the price level because it affects the relative supply of low and high skilled labor. Note also that 100% unemployment will not occur, because in this case no non-tradable goods would be produced.

[11] See Rogoff (1996) for a recent survey.

are considered in the simulations, the countries that are considered comprise about 93% of world population. Next, the table shows the mean per capita income, in PPP dollars, followed by the PPP deflator. The level of gross national product (GNP) per capita varies from a low of $414 (for Sierra Leone) to $38,247 (for Luxembourg). The next two columns present the Gini coefficients taken from World Bank (2000) or Deininger and Squire (1996), and the year for which the coefficient was calculated. There is significant variation in these coefficients, ranging as low as 0.19 for the Slovak Republic and exceeding 0.60 for Brazil, the Central African Republic, Gabon, Malawi and Sierra Leone. Although recall that the Gini coefficients are not precisely comparable, the wide range strongly suggests that inequality varies greatly across countries.

We assume that the utility function is Cobb–Douglas ($r = 0$). This leaves three world-wide parameters to be selected: δ, the share of tradables in consumption; d, the productivity parameter in the non-tradable sector; and α, the share of leisure. There are also two country-specific parameters: The extent of inequality and the average skill (a) level (the distribution of a is assumed to be log-normal). Finally, the calibration procedure requires that each country's revenue constraint is satisfied under the baseline tax system, adding the third country-specific requirement and pinning down the demogrant (G) under the baseline tax system.

In calibrating the model, we seek to match actual data regarding economy-specific mean incomes, Gini coefficients, and PPP indices, plus an overall world-wide average labor supply of 0.25. To achieve this, we first assume a standardized tax system with $t=0.3$ in all countries. Then, given d, δ and α, we adjust the distribution of skills in each country to exactly match the empirical mean income and relevant Gini coefficient. This requires solving for an equilibrium at each step, and yields the price of non-tradables and consumption of the two types of commodities. Having this information for all countries makes it possible to compute the PPP indices. [12]

The next step is to select the values of d, δ, and α that generate average labor supply at the desired level and that minimize the sum of squared deviations of the simulated PPP levels of 118 countries from their actual 1999 PPP price levels. [13] This

---

[12] We use the Eltetö–Köves–Szulc (EKS) method that was used to compute PPP in our data. See Hill (1997) for a discussion of purchasing power parity methods and the EKS formula (Eq. (50)).

[13] There are a few complications in implementing this method. Most importantly, it may not be possible to match the empirical Gini values even by choosing extreme values of the inequality of skills. To see this concern, consider the case when $\delta = 0$. In this situation, all individuals employed in the non-tradable sector have exactly the same wage rate and exactly the same income. Because a given fraction of income must be spent on the production of this sector, this requires a big enough fraction of population working in this sector. As the result, the combination of a relatively low value of δ and a relatively low value of d (i.e., a high share of non-tradables) makes the lower end of the distribution equal and large, therefore limiting the overall level of inequality. It turns out that there is a region of values of these parameters where the actual Gini coefficients for the most unequal countries may not be matched. It also turned out that the best choice of these parameters (i.e., the one that minimizes the deviations from the actual PPP levels) is on the boundary of this region (i.e., the country with the highest inequality level has an extreme inequality of skills). The parameters we use are almost on this boundary, but the results are not sensitive to shifting away from the boundary.

procedure generates calibrated parameter values of $\delta = 0.79$, $d = 0.12$, $\alpha = 0.63$. These parameter values imply that almost 80% of income is spent on tradable commodities. Furthermore, the small value of d implies that the non-tradable sector has significant decreasing returns to scale in individual skills, so that it is quite close to relying on just the amount of time provided. [14]

Table 1 presents the results of the calibration exercise for a few selected countries [15] as well as the world average. [16] The first column of Table 1 shows the average labor income if everybody worked full-time (i.e., consumed no leisure at all). The following columns show the average labor supply, consumption, and unemployment rate. Unemployment in this model is voluntary, and is a result of the demogrant that implies that a certain degree of consumption is possible even with zero labor supply. Although the simulated unemployment rate is as high as 21% for a few countries with large inequality (those with Gini coefficients exceeding 0.55), in aggregate only a tiny fraction (less than 0.5%) of the world's population chooses not to work. Those that choose to be unemployed are at the bottom of the ability distribution in a given country. Because with a Cobb–Douglas utility function, richer economies are just scaled-up versions of poorer ones, the unemployment rate is simply a function of the degree of inequality in underlying abilities.

The next two columns show the Gini coefficients of pre-tax labor income and consumption in the baseline simulation. Note that, because of the redistributive nature of the baseline tax system, the former is always higher than the latter, with the difference between the two measures ranging between 5 and 25 points. In each case, the parameters have been selected so that the relevant one of these is equal to the empirical value from Table 6. The Gini coefficient of consumption for the world as a whole is 0.68, while the Gini coefficient based on labor income is 0.72.

The final three columns show consumption levels at the 5th, 50th, and 95th percentiles of the distribution. Huge inequality of consumption is evident in the statistics for the world: Median consumption is $1814, while consumption at the 95th percentile is $27,860. As Figs. 1 and 2 show, the calibrated PPP indices quite closely match the actual ones, although this mostly reflects the fact that the dependence of the price level on income is well accounted for. In reality there is also significant variation in the price level conditional on income level, and this is not well explained by our model. There is a small variation of the price level conditional on income that is produced by our model (due to differences in inequality levels), but it is nowhere near what is observed in the data.

---

[14] With $d = 0$, $a^d = 1$, implying that skill would not matter at all in the non-tradable sector. As a result, only hours worked in that sector would determine its output.

[15] The appendix to the working paper version of this paper (Kopczuk et al., 2002) contains complete (i.e., for all 118 countries) versions of all the tables presented in this paper.

[16] The world average is computed over all individuals, and is not equal to the unweighted average of the country averages.

Table 1
Summary statistics about the baseline calibrated world economy, selected countries

|	Mean full time income	Mean labor supply	Mean consumption	Unemployment (%)	Pre-tax labor income Gini	Consumption (After-tax income) Gini	Consumption percentiles		
							5%	50%	95%
World	15,849	0.25	5060	0	0.72	0.68	609	1814	27,860
United States	95,093	0.27	30,636	0	0.41	0.29	17,692	21,206	70,711
Israel	55,230	0.27	17,458	0	0.36	0.25	10,817	13,372	36,518
Poland	12,639	0.28	3963	0	0.33	0.23	2538	3157	7896
Peru	7278	0.26	2391	0	0.46	0.32	1298	1508	5982
El Salvador	5635	0.25	1898	0	0.52	0.36	960	1142	5077
Papua New Guinea	2084	0.20	801	0	0.73	0.51	304	410	2143
India	1319	0.25	449	0	0.54	0.38	221	266	1230
Kyrgyz Republic	934	0.27	300	0	0.41	0.28	175	210	680
Ethiopia	289	0.24	100	0	0.57	0.40	47	58	285
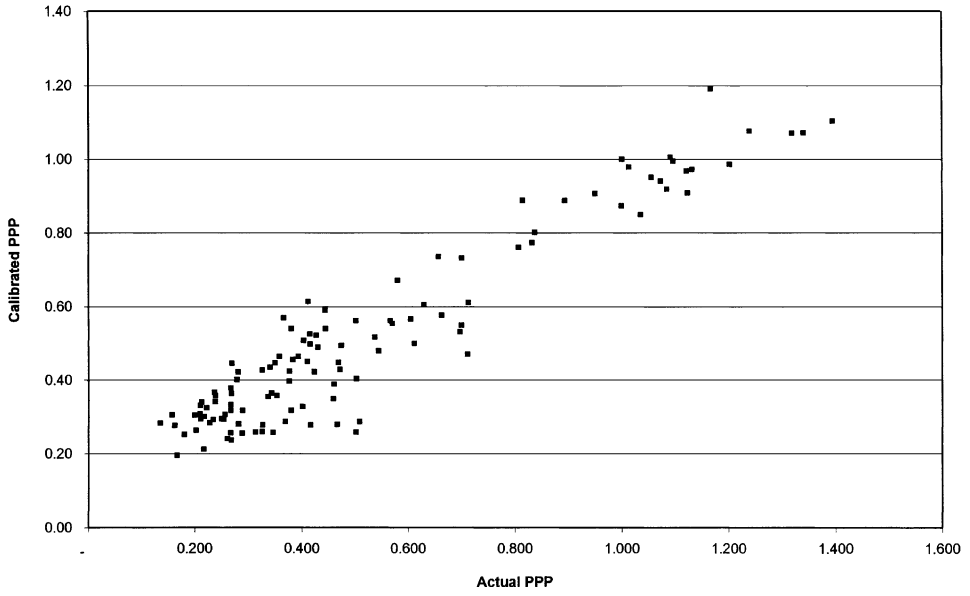
Fig. 1. Actual vs. calibrated PPP.
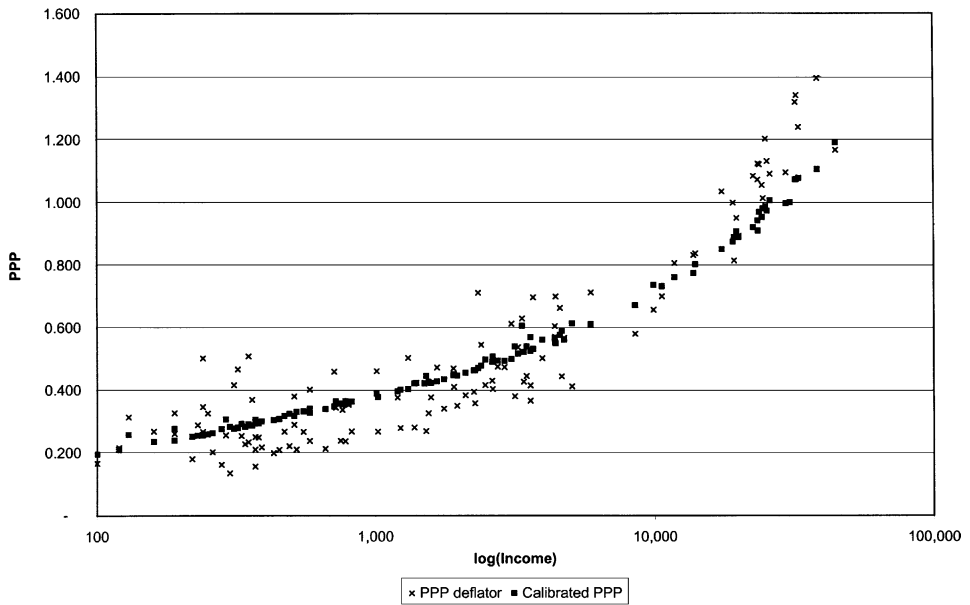
Fig. 2. Income vs. PPP.

3. Results

We are now ready to calculate the optimal income tax systems, first for each country choosing its own tax-and-transfer system, and then for the world income tax. Table 2 shows the results for a subset of countries.

In the focal simulation we assume that the parameter of the Atkinson's welfare function is $v = 2.0$.[17] The first and third columns of Table 2 show the parameters – marginal tax rate and demogrant – of the decentralized optimal linear income taxes. The optimal marginal tax rates are monotonically related to the Gini coefficients shown in Table 1. Under the decentralized solution, the optimal marginal tax rate varies between 0.13 for the Slovak Republic and 0.82 for Gabon. The population-weighted-average marginal tax rate is 0.41.

The second and fourth columns of Table 2 show the parameters of the optimal world income tax. The marginal tax rate is 0.62, substantially higher than the average under the decentralized solution, although smaller than the decentralized optimal tax rates for a handful of countries with very high inequality. The world income tax system also features a demogrant of $3112, which exceeds the actual per capita GNP for 73 countries. Note, however, that the aid from abroad backfires as well, because it takes the form of tradable commodities. As a result, the larger the aid, the lower the value of tradables in terms of non-tradables and the less effective is a dollar of transfers.

Because of the monotonic relationship between the Gini and optimal progressivity, the world income tax rate is higher than the rate for almost all countries in the world. For this reason, the deadweight loss is significantly higher than would occur under the decentralized system. This is, however, the cost of achieving a much larger amount of redistribution than would occur under the decentralized solution. This cost is justified by the better targeting of transfers compared to the decentralized solution: The centralized tax system does not provide demogrants (and therefore require higher-than-otherwise distorting marginal tax rates) to people who are poor from a country's perspective, but who are not poor from a world perspective.

The middle four columns of Table 2 show the Gini coefficients of consumption and labor income under the decentralized and world income tax regimes. Not surprisingly, the Gini coefficients of consumption (equal to after-tax income) are lower than those of labor income.[18] Redistributive tax systems render consumption considerably more equal.

A striking result of this simulation is that the decentralized tax system does not substantially affect the degree of inequality for the whole world. The Gini coefficient of consumption decreases only slightly when compared to the original calibrated world featured in Table 1.[19] In fact, if tax rates and demogrants in all countries were set

---

[17] As we discuss later, the qualitative conclusions are robust to changes in this parameter.

[18] In most cases the Gini coefficient of consumption falls below the baseline values of Table 1, for both the decentralized and the world tax systems, with exceptions to this rule being the economies that optimally set taxes below the baseline value of $t = 0.3$.

[19] This is possible because for many richer economies our baseline tax rate of 0.3 exceeds the optimal marginal tax rate, and therefore for these countries there is more redistribution in the baseline case than under the optimal income tax structure.

Table 2
Comparison of the decentralized solution and the World Income Tax

	Tax rate		Demogrant		Labor Gini[a]		Consump. Gini[a]		Mean labor supply		Mean consumption		Mean labor income		Unemployment		Per-capita transfer
	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	WIT
World	0.41	0.62	1539	3112	0.75	0.79	0.69	0.25	0.21	0.09	5027	5016	5027	5016	2%	15%	0.0
United States	0.36	0.62	10,373	3112	0.43	0.44	0.28	0.35	0.25	0.29	29,058	15,072	29,058	31,504	0%	0%	−16,432.0
Israel	0.30	0.62	5267	3112	0.36	0.41	0.25	0.27	0.27	0.25	17,417	9293	17,417	16,283	0%	0%	−6989.7
Poland	0.28	0.62	1127	3112	0.32	0.45	0.23	0.08	0.28	0.08	4041	3737	4041	1646	0%	0%	2090.2
Peru	0.40	0.62	876	3112	0.50	0.53	0.30	0.07	0.22	0.06	2179	3557	2179	1172	0%	9%	2384.3
El Salvador	0.45	0.62	751	3112	0.57	0.57	0.31	0.07	0.19	0.05	1660	3522	1660	1081	0%	15%	2441.2
Papua New Guinea	0.63	0.62	416	3112	0.82	0.72	0.31	0.08	0.08	0.04	665	3478	665	966	19%	40%	2512.3
India	0.47	0.62	182	3112	0.60	0.53	0.32	0.04	0.18	0.04	388	3384	388	718	0%	23%	2666.2
Kyrgyz Republic	0.35	0.62	101	3112	0.42	0.41	0.27	0.03	0.25	0.04	286	3381	286	710	0%	10%	2671.0
Ethiopia	0.50	0.62	42	3112	0.64	0.56	0.32	0.05	0.16	0.04	85	3381	85	710	0%	27%	2670.8

[a]Gini coefficients for the world are calculated using labor/consumption adjusted for purchasing power parity differences.

to zero, the Gini coefficient of consumption would be 0.695, compared to just 0.689 under the decentralized tax systems. Each country redistributing on its own makes only a small dent in world inequality. This result simply reflects that inequality in the distribution of all individuals' income, regardless of where in the world they live, is higher than the inequality of individuals' income within nearly every country of the world. According to Milanovic (2002), the differences in countries' mean income explain at least three-quarters of overall world inequality. No country on its own can transfer income from the world's rich to the world's poor, because no country has the world's poorest *and* the world's richest among its citizens. Consequently, decentralized redistribution cannot significantly address world's inequality.

The world income tax fares significantly better in reducing the inequality of consumption. The Gini coefficient goes from 0.69 under the decentralized tax regimes to 0.25 under the WIT, when calculated using consumption adjusted for the (endogenous) price level. However, because of its disincentive effects, the world income tax also decreases the average level of consumption and reduces average labor supply. Average labor supply (the number of hours worked) falls by more than half, from 0.21 to 0.09, under the world income tax. This decline is mostly due to the sharp decline in labor supply in the poor economies. The world unemployment rate increases from 2% to 15%.

Although by construction there are no cross-border transfers under the decentralized solution, under the world income tax the implicit transfers are substantial. For example, per capita the United States transfers $16,432 abroad. Countries at about the mean income of Uruguay and below receive net transfers, and the poorest countries receive more than $2600 per capita. The mean level of welfare for the whole world [20] increases under the world income tax system when compared to the decentralized solution, implying that the world income tax is more successful in redistributing income than the decentralized system. Under the decentralized solution, the average welfare level in the world is about equal to that of the average Filipino. A conversion to the world income tax brings it to about the level of a typical Czech. Not surprisingly, there are huge welfare gains for residents of the poor countries and substantial welfare losses for residents of the developed economies.

Fig. 3 illustrates the implications of switching from decentralized income tax systems to a world income tax by plotting tax as a function of gross income for individuals at the 5th, 50th, and 95th percentiles (under the decentralized system) for three countries: India, Poland, and the United States. For the citizens of the United States, tax due under the world income tax exceeds tax liability under the decentralized tax system for any level of gross income. This is also true for the richest Poles, but most Poles would observe a decrease in tax liability, absent behavioral response. Even the richest Indians gain, although not as much as the poorest ones. The figure also shows that the marginal tax rate increases under the WIT for all three economies. As a result, within each country the richest citizens gain least (or lose most).

The value of the substantial cross-country transfers may seem to be magnified by differences in the cost of living in different economies. For example, under the

---

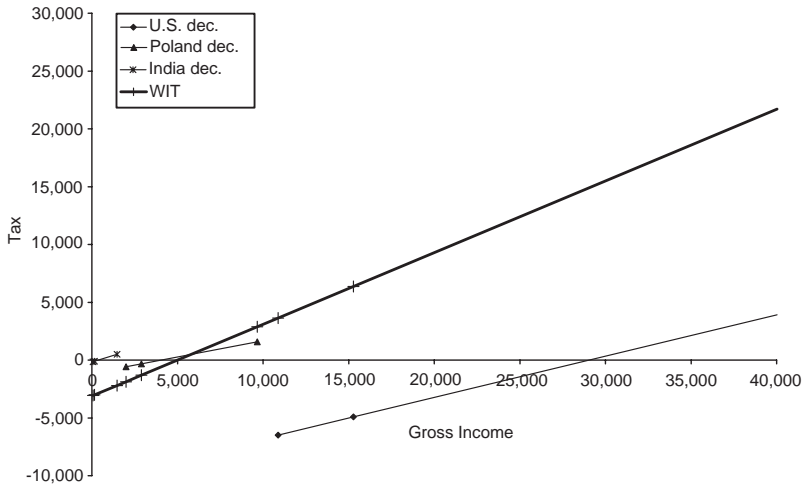[20] The welfare levels are normalized for expositional purposes. Only relative differences are of interest.

Fig. 3. Optimal tax functions.

decentralized solution the purchasing power of a dollar in India is magnified by a factor of more than three. However, transfers may take the form of tradables only, so that they are not as beneficial as a pure income transfer would be. This is reflected in significant changes in the cost of living of the poorest economies, which reflect the increased prices of non-tradables. This occurs because, as an economy becomes richer (due to transfers), the demand for non-tradables increases, but their supply is still bounded by the economy's own labor resources. In the case of India, the price of non-tradables under the WIT increases by a factor of twenty, and the overall cost of living increases from 0.33 to 0.78 (Table 3). In fact, one result of this transfer scheme is that most of the poorest economies end up consuming less non-tradable goods. This is because there is an overall decrease in labor supply as the result of the large transfer.

The differences in the cost of living also make it possible for average consumption in the world to stay almost constant in PPP terms. Looking at the percentiles of consumption, it is clear that under the world income tax most of the population gains. The consumption level of the world-median individual increases by $4600 in PPP terms. At the same time, the structure of consumption changes. Consumption of non-tradables falls in every country.

4. Foreign aid and the Bono/O'Neill factor

A striking feature of the hypothetical centralized scheme of world redistribution motivated by the pursuit of social justice, as represented by the optimal world income tax solution, is the large transfers from the rich countries, amounting in the United States to $16,432 per capita. In fact, many relatively well-off countries do provide

Table 3
Comparison of the decentralized solution and the WIT: Further details

| | Average non-tradable consumption | | Price of non-tradables | | PPP | | Consumption percentiles[a] | | | | | | Social welfare per capita (utils) | |
| | | | | | | | 5% | | 50% | | 95% | | | |
	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT
World	0.22	0.13	3371	7687			566	4962	1599	6198	28,784	14,043	−7,564,213	−4,384,246
United States	0.35	0.25	17,525	12,467	1.00	1.00	17,373	6835	20,198	10,808	65,683	37,657	−2,698,305	−3,608,738
Israel	0.34	0.21	10,858	9250	0.85	0.90	10,812	5260	13,346	7353	36,416	19,906	−3,180,572	−3,932,527
Poland	0.28	0.14	2992	5578	0.57	0.69	2566	3261	3216	3531	8097	5195	−4,961,715	−4,422,498
Peru	0.25	0.13	1869	5883	0.50	0.70	1272	3112	1465	3431	5207	4252	−5,846,696	−4,458,642
El Salvador	0.22	0.12	1587	6105	0.48	0.71	955	3112	1120	3399	4082	3945	−6,288,349	−4,474,412
Papua New Guinea	0.12	0.10	1198	7525	0.44	0.82	416	3112	480	3233	1251	4256	−8,109,220	−4,549,222
India	0.18	0.10	456	7100	0.33	0.78	222	3112	263	3327	963	3892	−9,720,412	−4,542,144
Kyrgyz Republic	0.20	0.10	299	7191	0.29	0.79	172	3112	201	3362	635	3734	−10,884,054	−4,549,296
Ethiopia	0.14	0.08	127	8589	0.22	0.91	49	3112	58	3311	212	3934	−15,303,253	−4,609,515

[a]Consumption percentiles for the world are calculated using consumption adjusted for purchasing parity differences.

foreign aid to less well-off countries, and most rich countries contribute to multilateral institutions such as the World Bank that offer assistance to relatively poor countries. How does it compare to our simulated level of optimal transfers, and what does the comparison imply?

4.1. Foreign aid

The Development Assistance Committee (DAC) of the OECD publishes annual data on both bilateral and multilateral aid flows.[21] Its 1999 report indicates that in 1998 the U.S. gave $5.988 billion of bilateral assistance, and $2.798 billion of multilateral assistance, for a total of $8.786 billion of official development assistance. This last figure represents 0.10% of U.S. GNP, and is $33 per capita. To put the U.S. figures in perspective, for all 21 DAC countries (including the U.S.), official development assistance represented 0.24% of GNP; the U.S. ranks 21st among the 21 countries represented. The actual amount of net aid contributed or received by various countries, from World Bank (2000, Tables 6.8 and 6.10), is presented in the last column of Table 6.[22]

4.2. Bono and O'Neill: Estimating the implicit discounting of foreigners' well-being and/or the implicit extent of waste of the decentralized outcome

The actual flows of foreign aid are miniscule compared to what our simulated world income tax generates. We have demonstrated that the discrepancy cannot be explained by the efficiency costs that would result from the higher marginal tax rates needed to generate the tax revenue to be transferred from the poor countries – that is an integral part of the WIT simulations. One natural explanation for the discrepancy is that, contrary to the assumption of the model of a centralized decision-maker motivated by the pursuit of border-neutral social justice, the residents of rich countries such as the United States are not border-neutral at all, but rather value the welfare of a foreigner

---

[21] There is a considerable literature on the determinants of foreign aid, in particular the extent to which it is motivated by strategic and political considerations as opposed to altruistic and humane ones. Lumsdaine (1997) investigates the effect of colonial links between donor and recipient, the democratic status of the recipients, and the income level of the recipient, but presents only simple correlations rather than a full-blown multivariate analysis. Alesina and Dollar (2000) do perform such an analysis (of bilateral aid flows only), and find considerable evidence that the direction of foreign aid is indeed dictated by political and strategic considerations much more than by either the economic needs or the policy performance of the recipient. A separate but relevant literature concerns the effects of foreign aid on the receiving countries, and has been studied by Jepma (1997) and Boone (1994, 1996). Most recently, Burnside and Dollar (2000) find that aid is beneficial to countries that adopt appropriate and stable policies, and is wasted otherwise. However, they find no evidence that foreign aid encourages the adoption of "good" macroeconomic policies.

[22] Table 6.8 of World Bank (2000) reveals the official development assistance and aid contributions of the high-income economies in 1998. It includes both bilateral transfers and contributions to the financial institutions. Table 6.10 shows the amount of assistance and aid received by various countries. These numbers do not balance out. This is because some aid is allocated by region, but not by country, and because of administrative costs, research into development issues, and aid to non-governmental organizations. As the result, contributions exceed aid received by approximately $22 billion. The total amount of aid received is about $35 billion.

significantly less than the welfare of their own residents with the same income. A second is that transfers are not used efficiently, so that the richer countries perceive them as a waste of resources. Finally, part of the discrepancy could be caused by the failure of the decentralized system to overcome the inherent free-rider problem faced by altruistic donors.

The notion that the sense of justice of Americans (or residents of other countries) stops, or nearly stops, at the border will not shock most readers. Neither will the possibility that transfers are wasted. With the model we have developed, though, we can go beyond suggesting these notions to quantify what the actual flows of aid imply about how much the United States implicitly weighs the well-being of a resident of, say, India. Our weights reflect a combination of a lower weight put on foreigners' well-being and the extent of waste.

To fix ideas, consider a simple version of this setup in which the U.S. and India are the only countries in the world, and each country has only a poor person (denoted P) and a rich person (denoted R).[23] Each country makes its own decisions about its tax-and-transfer system to maximize a Bergson–Samuelson formulation of social justice constrained by the disincentives effects of redistributive tax-and-transfer systems. The social welfare function of the U.S. includes the utility level of Indians, although the Indians' utilities may have a relative weight of less than one. The social welfare function of the U.S. has the form $W = (1-v)^{-1}[U_{RS}^{1-v} + U_{PS}^{1-v} + bU_{RI}^{1-v} + bU_{PI}^{1-v}]$, where U_{ij} refers to the utility of the ith person in the jth country (S = US and I = India), and b $(0 < b)$ is the relative weight placed on an Indian's utility.

The United States now has three policy instruments: The demogrant and income tax rate as before, plus a transfer to the Indian government. The U.S. knows the Indian social welfare function, so it knows exactly how India will adjust its own demogrant and tax rate upon receipt of a transfer, and can therefore calculate the increase in the utility of each Indian citizen. Given these assumptions, we can in the framework of our simulated model calculate the amount of transfer to India the U.S. will make for any value of b. Conversely, we can work backwards and calculate what value of b is consistent with the amount of transfers we observe. In what follows we do the latter. Before we do so, we introduce the possibility that transfers are wasted. Specifically, we assume that a given fraction a^i of the transfer is wasted by corrupt politicians, whose welfare we assign a zero weight.

Calculating the implied weights on the well-being of other countries' residents is straightforward, as long as each country selects its tax system optimally. Denote by λ^i the marginal social welfare benefit from a marginal increase in public resources in country i. Formally, this is the Lagrange multiplier on the revenue constraint in the ith country's optimal tax problem.[24] Because of the possibility of waste, the marginal welfare from the transfer of a dollar to country i is then $(1 - a^i)\lambda^i$. On the margin, the optimizing government considering international aid compares its own λ to that

---

[23] By considering only one rich country we put aside the possibility that altruistic impulses are thwarted by the free-rider problem.

[24] At the optimum, it is equal to the average of marginal utilities of income (from the social welfare point of view) in a given country.

Table 4
Implied U.S. weights

	Decentralized solution	WIT
World	0.1591	0.3795
United States	1.0000	1.0000
France	0.8188	0.9051
Israel	0.5284	0.6618
Poland	0.0802	0.2933
Peru	0.0336	0.2771
El Salvador	0.0233	0.2725
Papua New Guinea	0.0071	0.2524
India	0.0035	0.2609
Kyrgyz Republic	0.0024	0.2636
Ethiopia	0.0005	0.2557

of other countries. At the optimum, the government of donor country i must then set $\lambda^i = (1 - a^j) b^j \lambda^j$, for any recipient country j, where b^j is the welfare weight attached to country j. This formula allows us to calculate the product $(1 - a^j) b^j$ directly, because optimization yields the values of the λ's.

Table 4 presents the implied marginal weights from the point of view of the U.S. for a selected group of countries. By construction, ceteris paribus, weights for the poorer economies must be smaller than those for the richer ones. For the poorest economy of Ethiopia, this weight is just 0.0005. One blunt interpretation is that the latter number implies that the amount of actual foreign aid given by the U.S. to Ethiopia is consistent with the well-being of an Ethiopia resident being valued at 1/2000 of that of an American. Alternatively, it can be believed that only 1/20th of one percent of aid reaches its desired recipients. A combination of the two is also possible. For example, if as much as 5% of aid reaches its recipients, the corresponding welfare weight consistent with the observed amount of aid would still be equal to just 0.01.

Implied welfare weights from the point of view of any other country can be derived from the Table 4 by re-normalizing the countries' multiplier using the Lagrange multiplier of an alternative donor rather than the U.S. For example, welfare weights from the point of view of France can be obtained by dividing the first column of Table 4 by 0.8188 and the second one by 0.9051. Similarly, welfare weights from the point of view of Norway can be obtained by dividing numbers in the first column by 1.4493 and numbers in the second column by 1.4359 – these are weights of Norway relative to the U.S. Despite the larger level of per-capita aid provided by Norway than by the U.S., Norway remains wealthier than the U.S. and it should still find it even more attractive than the U.S. to provide aid to the poor countries.

The column labeled WIT in Table 4 reveals that even under an optimal world income tax there is still room for a potential welfare improvement: The average weight for the rest of the world is 0.4, so that a marginal dollar in U.S. transfers would still finance a $2.50 increase in welfare, if used to finance a universal increase in the demogrant.

This is, however, not feasible in our model because of the assumed linearity of the tax system that precludes a unilateral change of the U.S. transfers. [25]

5. Sensitivity analyses

In Table 5, we present the results of simulations analogous to those of Section 3, but for different degrees of concavity of the common social welfare function. We consider $v = 0.5$, 2.0 and 5.0. Because $v = 2.0$ is our baseline case, the numbers in this part repeat information shown earlier. A value of $v = 5.0$ corresponds to a much more egalitarian social welfare function, while $v = 0.5$ is a much less egalitarian social welfare function. To save space, we show only the results for the world as a whole and three different countries: The United States, Poland and India. As expected, increasing egalitarianism leads to more redistribution: Marginal tax rates increase under both the decentralized and world income tax solutions. Notably, though, the changes are much larger in the decentralized case. This is because world inequality is very extreme to begin with, and therefore even a low redistributive incentive induces high marginal tax rates (and the optimal marginal tax rate is bounded from above by the one corresponding to the "peak" of the Laffer curve). Indeed, the optimal world income tax is almost unaffected by changes in the concavity of the welfare function.

Changes in the social welfare function also have significant consequences for the implied weights. This is intuitive. Without any redistributive incentive, these weights would all be equal to one even if the distribution of incomes were very unequal. Therefore, the lower is the concavity of the social welfare function, the higher should be these weights. For example, when $v = 0.5$, the implied welfare weight attached by the U.S. to an Indian is 0.0235, and it falls to 0.0001 for $v = 5.0$.

6. Summary and ruminations

The decentralization of redistribution decisions results in vastly less redistribution than would be accomplished by a centralized scheme motivated by the pursuit of border-neutral social justice taking into account the disincentive effects caused by the higher taxes needed for cross-country transfers. In our stylized simulation of redistribution policy, a decentralized system without cross-country transfers hardly budges the world Gini coefficient of consumption, even though it reduces it for particular countries. Put bluntly, within-country redistributive schemes are of almost no value from the world perspective. In contrast, a world income tax would provide a drastic reduction in consumption inequality, cutting the Gini coefficient by nearly two-thirds. The decentralized scheme is also relatively inefficient, as it causes an efficiency loss that is larger than it need be to achieve the same amount of redistribution as would a centralized system. To be sure, the world income tax features a much higher absolute efficiency

---

[25] In a more general nonlinear tax system the feasibility of such transfers would be limited by the incentive constraints.

Table 5
Sensitivity analysis

| | Marginal tax | | Demogrant | | Transfer | PPP | | Consumption percentiles | | | | | | U.S. welfare weights | |
| | | | | | | | | 5% | | 50% | | 95% | | | |
	Dec.	WIT	Dec.	WIT	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT	Dec.	WIT
$v = 0.5$															
World	0.36	0.60	1329	3061	0.0			574	5046	1677	6292	29,970	14,595	0.1982	0.4266
United States	0.30	0.60	9244	3061	−16,155.7	1.00	1.00	17,691	7077	21,167	11,203	70,509	39,202	1.0000	1.0000
Poland	0.21	0.60	913	3061	2010.4	0.57	0.68	2641	3243	3390	3513	8714	5349	0.2018	0.3506
India	0.43	0.60	172	3061	2632.5	0.33	0.76	223	3061	265	3288	1031	3848	0.0235	0.3201
$v = 2.0$															
World	0.41	0.62	1539	3112	0.0			566	4962	1599	6198	28,784	14,043	0.1591	0.3795
United States	0.36	0.62	10,373	3112	−16,432.0	1.00	1.00	17,373	6835	20,198	10,808	65,683	37,657	1.0000	1.0000
Poland	0.28	0.62	1127	3112	2090.2	0.57	0.69	2566	3261	3216	3531	8097	5195	0.0802	0.2933
India	0.47	0.62	182	3112	2666.2	0.33	0.78	222	3112	263	3327	963	3892	0.0035	0.2609
$v = 5.0$															
World	0.45	0.64	1722	3162	0.0			561	4864	1527	6081	27,408	13,433	0.1385	0.3273
United States	0.41	0.64	11,252	3162	−16,713.8	1.00	1.00	17,014	6570	19,233	10,375	61,132	35,966	1.0000	1.0000
Poland	0.33	0.64	1280	3162	2174.2	0.57	0.70	2493	3277	3060	3548	7572	5024	0.0129	0.2245
India	0.51	0.64	189	3162	2699.5	0.34	0.80	221	3162	261	3363	905	3933	0.0001	0.1883

cost, because it has a higher marginal tax rate than most countries would choose on their own.

In reality there are cross-country transfers in the form of foreign aid. However, the actual flow of foreign aid is minuscule compared to what the optimal world income tax implies, suggesting that voluntary world redistribution produces an outcome that is consistent with the residents of rich countries not being border-neutral, or anything close to it, or with assuming that most of the aid will be wasted. In our baseline case, we calculate that the current situation is consistent with the U.S. on average valuing the well-being of foreigners only 1/6 as much as an American citizen, and less than 1/2000 for the residents of the poorest of the developing economies. Alternatively, this latter value corresponds to an extreme extent of waste so that only 1/20 of one percent of transfers reaches its desired recipients.

This conclusion is sensitive to the assumed concavity of the social welfare function. Furthermore, our interpretation of weights is subject to a number of caveats. The first is due to the restrictiveness of the instruments that we consider: A linear tax does not allow the targeting of aid directly to the poorest members of the poor economies. If more targeted ways of transferring aid were available, the implied weights consistent with actual transfers would be even lower. We consider only a static framework and do not account for the effect that transfers can have on human or physical capital accumulation and, therefore, on future growth. Finally, it would certainly be interesting to disentangle the failure of the existing decentralized system to approach the centralized redistribution solution motivated by the pursuit of social justice into its three components – lack of border neutrality in the consideration of social justice, the perceived wastage of foreign aid, and the inability of a decentralized system to overcome the free-rider problem that limits the expression of the altruistic impulses of rich countries.

Appendix

Table 6 lists the key data for all of the 118 countries.

Table 6
Data

	Country	Population (in 1000)	GNP per capita	GNP per capita (PPP)	PPP deflator	Gini coefficient	Year of Gini	Net aid per capita
1	Luxembourg	432	44,640	38,247	1.167	0.27	1994[g]	−266.20
2	Switzerland	7120	38,350	27,486	1.395	0.33	1992[g]	−136.80
3	Norway	4454	32,880	26,522	1.240	0.26	1995[g]	−308.26
4	Japan	126,570	32,230	24,041	1.341	0.25	1993[g]	−85.11
5	Denmark	5317	32,030	24,280	1.319	0.25	1992[g]	−342.67
6	United States	272,878	30,600	30,600	1.000	0.41	1997[g]	−42.19
7	Singapore	3223	29,610	27,024	1.096	0.39	1989[g,d]	0.62
8	Austria	8086	25,970	23,808	1.091	0.23	1987[g]	−80.01
9	Germany	82,027	25,350	22,404	1.131	0.30	1994[g]	−76.01

Table 6 (*continued*)

	Country	Population (in 1000)	GNP per capita	GNP per capita (PPP)	PPP deflator	Gini coefficient	Year of Gini	Net aid per capita
10	Sweden	8857	25,040	20,824	1.202	0.25	1992[g]	−189.45
11	Belgium	10,223	24,510	24,200	1.013	0.25	1992[g]	−93.03
12	Netherlands	15,802	24,320	23,052	1.055	0.33	1994[g]	−200.73
13	Finland	5167	23,780	21,209	1.121	0.26	1991[g]	−92.51
14	Hong Kong	6877	23,520	20,939	1.123	0.45	1991[g,d]	1.02
15	France	60,794	23,480	21,897	1.072	0.33	1995[g]	−107.99
16	United Kingdom	59,110	22,640	20,883	1.084	0.36	1991[g]	−72.73
17	Australia	18,994	20,050	22,448	0.893	0.35	1994[g]	−50.59
18	Italy	57,649	19,710	20,751	0.950	0.27	1995[g]	−43.73
19	Canada	30,604	19,320	23,725	0.814	0.32	1994[g]	−60.38
20	Ireland	3727	19,160	19,180	0.999	0.36	1987[g]	−53.39
21	Israel	6093	17,450	16,867	1.035	0.36	1992[g]	174.95
22	Spain	39,410	14,000	16,730	0.837	0.33	1990[g]	−35.04
23	New Zealand	3823	13,780	16,566	0.832	0.44	1993[g]	−34.00
24	Greece	10,536	11,770	14,595	0.806	0.33	1993[g]	−16.99
25	Portugal	9990	10,600	15,147	0.700	0.36	1994–95[g]	−27.93
26	Slovenia	1981	9890	15,062	0.657	0.27	1995[g]	20.19
27	Korea Republic	46,848	8490	14,637	0.580	0.32	1993[c]	−1.07
28	Uruguay	3312	5900	8280	0.713	0.42	1989[g]	7.25
29	Czech Republic	10,280	5060	12,289	0.412	0.25	1996[g]	43.48
30	Chile	15,018	4740	8370	0.566	0.57	1994[g]	6.99
31	Hungary	10,068	4650	10,479	0.444	0.31	1996[g]	20.76
32	Croatia	4464	4580	6915	0.662	0.27	1998[c]	8.74
33	Brazil	168,066	4420	6317	0.700	0.60	1996[g]	1.96
34	Mexico	97,425	4400	7719	0.570	0.54	1995[g]	0.15
35	Trinidad and Tobago	1293	4390	7262	0.605	0.40	1992[g]	10.83
36	Poland	38,695	3960	7894	0.502	0.33	1996[g]	23.31
37	Venezuela	23,707	3670	5268	0.697	0.49	1996[g]	1.56
38	Slovak Republic	5396	3590	9811	0.366	0.20	1992[g]	28.72
39	Mauritius	1170	3590	8652	0.415	0.37	1991[c,d]	34.19
40	Estonia	1442	3480	7826	0.445	0.35	1995[g]	62.41
41	Malaysia	22,710	3400	7963	0.427	0.49	1995[g]	8.89
42	Gabon	1208	3350	5325	0.629	0.63	1977[c,d]	37.25
43	Botswana	1588	3240	6032	0.537	0.54	1986[c,d]	66.75
44	South Africa	21,429	3160	8318	0.380	0.59	1993–94[c]	23.89
45	Panama	2808	3070	5016	0.612	0.49	1997[c]	7.83
46	Turkey	64,328	2900	6126	0.473	0.42	1994[c]	0.22
47	Costa Rica	3588	2740	5770	0.475	0.47	1996[g]	7.53
48	Belarus	10,208	2630	6518	0.403	0.22	1998[c]	2.74
49	Lithuania	3699	2620	6093	0.430	0.32	1996[c]	34.60
50	Latvia	2430	2470	5938	0.416	0.32	1998[g]	39.92
51	Peru	25,230	2390	4387	0.545	0.46	1996[g]	19.86
52	Jamaica	2598	2330	3276	0.711	0.36	1996[c]	6.93
53	Russian Federation	146,512	2270	6339	0.358	0.49	1998[c]	6.94
54	Colombia	41,539	2250	5709	0.394	0.57	1996[g]	4.00
55	Tunisia	9457	2100	5478	0.383	0.40	1990[c]	15.65
56	Thailand	61,691	1960	5599	0.350	0.41	1998[c]	11.18

Table 6 (*continued*)

	Country	Population (in 1000)	GNP per capita	GNP per capita (PPP)	PPP deflator	Gini coefficient	Year of Gini	Net aid per capita
57	Dominican Republic	8404	1910	4653	0.410	0.49	1996[g]	14.28
58	El Salvador	6189	1900	4048	0.469	0.52	1996[g]	29.08
59	Iran	62,977	1760	5163	0.341	0.43	1984[c]	2.60
60	Guatemala	11,086	1660	3517	0.472	0.60	1989[g]	21.02
61	Paraguay	5359	1580	4193	0.377	0.59	1995[g]	14.18
62	Algeria	29,950	1550	4753	0.326	0.35	1995[c]	12.99
63	Romania	22,458	1520	5647	0.269	0.28	1994[g]	15.85
64	Jordan	4693	1500	3542	0.423	0.36	1997[c]	86.94
65	Egypt	62,430	1400	3303	0.424	0.29	1995[c]	30.67
66	Bulgaria	8216	1380	4914	0.281	0.28	1995[c]	28.24
67	Ecuador	12,409	1310	2605	0.503	0.44	1995[c]	14.18
68	Kazakhstan	15,438	1230	4408	0.279	0.35	1996[c]	13.41
69	Morocco	28,238	1200	3190	0.376	0.40	1998–99[c]	18.70
70	Philippines	76,785	1020	3815	0.267	0.46	1997[c]	7.91
71	Bolivia	8135	1010	2193	0.461	0.42	1990[g]	77.20
72	Sri Lanka	18,985	820	3056	0.268	0.34	1995[c]	25.81
73	Papua New Guinea	4705	800	2263	0.354	0.51	1996[c]	76.73
74	China	1,249,671	780	3291	0.237	0.40	1998[g]	1.89
75	Honduras	6325	760	2254	0.337	0.54	1996[g]	50.28
76	Ukraine	49,908	750	3142	0.239	0.33	1996[c]	7.61
77	Uzbekistan	24,600	720	2092	0.344	0.33	1993[g]	5.85
78	Cote d'Ivoire	14,729	710	1546	0.459	0.37	1995[c]	54.18
79	Turkmenistan	4779	660	3099	0.213	0.41	1998[c]	3.56
80	Cameroon	14,691	580	1444	0.402	0.49	1983[c,d]	28.86
81	Indonesia	207,022	580	2439	0.238	0.37	1996[g]	6.08
82	Lesotho	2105	550	2058	0.267	0.56	1986–87[c]	31.35
83	Zimbabwe	11,904	520	2470	0.211	0.57	1990–91[c]	23.52
84	Guinea	7247	510	1761	0.290	0.40	1994[c]	49.54
85	Senegal	9285	510	1341	0.380	0.41	1995[c]	54.07
86	Armenia	3809	490	2210	0.222	0.39	1989[g,d]	36.23
87	Pakistan	134,790	470	1757	0.268	0.31	1996–97[c]	7.79
88	India	997,515	450	2,149	0.209	0.38	1997[c]	1.60
89	Nicaragua	4919	430	2154	0.200	0.50	1991[c]	114.25
90	Ghana	18,949	390	1793	0.218	0.33	1997[c]	36.99
91	Mauritania	2598	380	1522	0.250	0.39	1995[c]	65.82
92	Vietnam	77,515	370	1755	0.211	0.36	1998[c]	15.00
93	Bangladesh	127,669	370	1475	0.251	0.34	1995–96[c]	9.80
94	Moldova	4281	370	2358	0.157	0.34	1992[g]	7.71
95	Kenya	29,410	360	975	0.369	0.45	1994[c]	16.12
96	Yemen	17,048	350	688	0.509	0.40	1992[c]	18.18
97	Mongolia	2623	350	496	0.234	0.33	1995[c]	77.39
98	Gambia	1251	340	1492	0.228	0.48	1992[c]	30.38
99	Sudan	28,993	330	1298	0.254	0.39	1968[g,d]	7.21
100	Uganda	21,479	320	1136	0.282	0.39	1992–93[c]	21.93
101	Zambia	9881	320	686	0.466	0.50	1996[c]	35.32

Table 6 (*continued*)

	Country	Population (in 1000)	GNP per capita	GNP per capita (PPP)	PPP deflator	Gini coefficient	Year of Gini	Net aid per capita
102	Nigeria	123,897	310	744	0.417	0.51	1996–97[c]	1.65
103	Kyrgyz Republic	4744	300	2223	0.135	0.41	1997[g]	45.53
104	Central African Republic	3540	290	1131	0.256	0.61	1993[c]	33.90
105	Lao PDR	5097	280	1726	0.162	0.30	1992[c]	55.13
106	Cambodia	11,757	260	1286	0.202	0.40	1997[c]	28.66
107	Madagascar	15,051	250	766	0.326	0.46	1993[c]	32.82
108	Tanzania	32,923	240	478	0.502	0.38	1993[c]	30.31
109	Mali	10,911	240	693	0.346	0.51	1994[c]	31.99
110	Burkina Faso	10,996	240	898	0.267	0.48	1994[c]	36.10
111	Mozambique	17,264	230	797	0.289	0.40	1996–97[c]	60.18
112	Nepal	23,384	220	1219	0.180	0.37	1995–96[c]	17.28
113	Malawi	10,788	190	581	0.327	0.62	1993[c,d]	40.23
114	Niger	10,493	190	727	0.261	0.51	1995[c]	27.73
115	Guinea-Bissau	1185	160	595	0.269	0.56	1991[c]	81.01
116	Sierra Leone	4949	130	414	0.314	0.63	1989[c]	21.42
117	Burundi	6678	120	553	0.217	0.33	1992[c]	11.53
118	Ethiopia	62,782	100	599	0.167	0.400	1991[c]	10.32

[c] Gini coefficient based on consumption or net income data.
[d] Value of Gini from Deininger and Squire (1996).
[g] Gini coefficient based on gross income data.

References

Alesina, A., Dollar, D., 2000. Who gives Foreign aid to whom and why? Journal of Economic Growth 5, 33–63.

Atkinson, A., 1970. The measurement of inequality. Journal of Economic Theory 2, 244–263.

Auten, G., Carroll, R., 1999. The effect of income taxes on household behavior. Review of Economics and Statistics 81, 681–693.

Balasa, B., 1964. The purchasing power parity doctrine: A reappraisal. Journal of Political Economy 76, 584–596.

Blundell, R., MaCurdy, T.E., 1999. Labor supply: A review of alternative approaches. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, Vol. 3. Elsevier/North-Holland, Amsterdam, New York.

Boone, P., 1994. The impact of foreign aid on savings and growth. Mimeo., London School of Economics.

Boone, P., 1996. Politics and the effectiveness of foreign aid. European Economic Review 40, 289–329.

Burnside, C., Dollar, D., 2000. Aid, policies, and growth. American Economic Review 90, 847–868.

Cooter, R., Helpman, E., 1976. Optimal income taxation and transfer payments under different social welfare criteria. Quarterly Journal of Economics 88, 656–670.

Deininger, K., Squire, L., 1996. A new data set measuring income inequality. World Bank Economic Review 10, 565–591.

Gruber, J., Saez, E., 2002. The elasticity of taxable income: Evidence and implications. Journal of Public Economics 84 (1), 1–32.

Helpman, E., Sadka, E., 1978. The optimum income tax: Some comparative static results. Journal of Public Economics 9, 383–393.

Hill, R.J., 1997. A taxonomy of multilateral methods for making international comparisons of prices and quantities. Review of Income and Wealth 43, 49–69.

Jepma, C.J., 1997. On the Effectiveness of Development Aid. World Bank, Washington, DC, unpublished.

King, R.G., Plosser, C.I., Rebelo, S.T., 1982. Production, growth and business cycles: I. The basic neoclassical model. Journal of Monetary Economics 21, 195–232.

Kopczuk, W., Slemrod, J., Yitzhaki, S., 2002. Why world redistribution fails. National Bureau of Economic Research Working Paper 9186.

Lumsdaine, D.H., 1997. Moral Vision in International Politics. Princeton University Press, Princeton, NJ.

Lydall, H.F., 1968. The Structure of Earnings. Clarendon Press, Oxford.

Milanovic, B., 2002. True world income distribution, 1988 and 1993: First calculation based on household surveys alone. Economic Journal 112, 51–92.

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

Rogoff, K., 1996. The purchasing power parity puzzle. Journal of Economic Literature 72, 647–668.

Sachs, J.D., 2001. One-tenth of 1 percent to make the world safer. Washington Post, November 21, p. A23.

Samuelson, P.A., 1964. Theoretical notes on trade problems. Review of Economics and Statistics 46, 145–154.

Slemrod, J., 1998. Methodological issues in measuring and interpreting taxable income elasticities. National Tax Journal 51, 773–788.

Stern, N.H., 1976. On the specification of models of optimum income taxation. Journal of Public Economics 6, 123–162.

World Bank, 2000. World Development Indicators 2000. World Bank, Washington, DC.