Risk & Sustainable Management Group

The Poverty Burden: A Measure of the Difficulty of Ending Extreme Poverty

John Quiggin[†]
&
Renuka Mahadevan

[†]Australian Research Council Federation Fellow, University of Queensland

THE UNIVERSITY
OF QUEENSLAND
AUSTRALIA

This version: 30 March 2010

The Poverty Burden: A Measure of the Difficulty of Ending Extreme Poverty

John Quiggin and Renuka Mahadevan*

School of Economics

University of Queensland

EMAIL j.quiggin@uq.edu.au

PHONE + 61 7 3346 9646

FAX   +61 7 3365 7299

http://www.uq.edu.au/economics/johnquiggin

Abstract

In this paper, we consider an extension of the commonly used poverty gap measure that may be used to address he question of whether the elimination of poverty is a feasible objective, given sufficient political commitment. The resources potentially available to address poverty may be measured by the total amount by which the incomes of the non-poor exceed the same poverty line. The ratio of the poverty gap to the resources potentially available is equal to the proportional tax rate on incomes in excess of the poverty line that would be required to fund a transfer sufficient to raise the incomes of all poor people to the poverty line. We refer to this ratio as the 'poverty burden' (PB). We provide a formal definition of the poverty burden and an analysis of its properties as a poverty measure.

The Poverty Burden: A Measure of the Difficulty of Ending Extreme Poverty

1. Introduction

For perhaps the first time in the history of the world, the majority of people alive today are not poor, or, at least not extremely poor. The World Bank (2008) estimates that 2.5 billion people live on incomes of less than $US2/day (in PPP-adjusted 1990 values), of whom nearly 1 billion live on less than $1/day. By contrast, nearly 4 billion people have incomes above the $US2/day poverty line. Of these, around 1 billion people are very rich, at least by historical standards, with incomes greater than $10 000 a year, or around $300 a day.

Relatively modest transfers from the well-off to the poor could, in the absence of leakages or offsetting negative effects, raise everyone in the world above the $1/day or $US2/day poverty lines. Thus, it is possible to propose the end of extreme poverty as a policy objective for the world community (Sachs 2005). The task has been made more difficult by the global financial crisis that began in 2008, but remains urgent.

The suggestion of eliminating poverty raises new issues of measurement. Although a wide range of poverty measures has been proposed, none seems well-suited to the question of whether the elimination of poverty is a feasible objective, given sufficient political commitment.

In this paper, we consider an extension of the commonly used poverty gap measure that may be used to address this question. The poverty gap is equal to the total amount by which the incomes of the poor fall short of a given poverty line. Conversely, the resources potentially available to address poverty may be measured by the total amount by which the incomes of the non-poor exceed the same poverty line. The ratio of the poverty gap to the resources potentially available is equal to the proportional tax rate on incomes in excess of the poverty line that would be required to fund a transfer sufficient to raise the incomes of

all poor people to the poverty line. We refer to this ratio as the 'poverty burden' (PB).

Measures of this kind are not entirely new. Anand (1977, p. 7) notes parenthetically, and without citation, that 'In the United States, the poverty gap is sometimes expressed as a fraction of GNP', and estimates the poverty gap for Malaysia at 7.3 per cent of total income, or 8.3 per cent of the income of the non-poor. Beckerman (1979) computes the poverty gap for the United Kingdom as a percentage of GDP and notes (p. 269) that 'the post-benefit gap is only 0.25% of G.D.P.' However, these measures have not been widely used, and there appears to have been little or no study of their formal properties.

The aim of this paper is to investigate the properties of the poverty burden, and to consider its usefulness in informing the debate over policies to reduce or eliminate extreme poverty. The paper is organized as follows:

Section 2 contains background material on measures of poverty and the way in which they have been used. Section 3 provides a formal definition of the poverty burden and an analysis of its properties as a poverty measure. Evidence on the way in which the poverty-income ratio has changed over time, for the world as a whole and for various regions is presented in Section 4. Section 5 deals with the relationship between the poverty burden measure and the policy debate over the effectiveness of aid in reducing poverty. It is argued that this debate has been distorted by an excessive focus on aid as an investment aimed at increasing rates of economic growth. Finally some concluding comments are offered in Section 6.

2. Background

A wide variety of poverty measures have been considered (Foster 1984, 1988; Kakwani 1980; Ravallion 1996; Sen 1976; Shorrocks 1987, 1995; Vaughan 1987). Most measures are computed with reference to a poverty line, which may be determined either exogenously, on the basis of some estimate of the expenditure required to avoid poverty, or endogenously, as a proportion of some measure of mean or median income. This difference in the choice of poverty line corresponds

to the distinction between absolute and relative measures of poverty (Foster 1998; Zheng 1994).

Prominent examples of poverty measures include the head count measure (the number of people, or proportion of the population, falling below the poverty line), the poverty gap (the aggregate amount by which incomes of the poor fall below the poverty line) and the Sen normalized poverty value (Sen 1976).

Broadly speaking, these measures are inequality measures that are sensitive only to the incomes of the poor. Changes in incomes for the non-poor affect these measures only to the extent that they may change the value of the poverty line, or to the extent that previously non-poor people become poor.

If the poverty line is determined exogenously, an increase in the income of a non-poor person has no effect on measured poverty. If the poverty line is (weakly) monotonically increasing in the income vector (for example, if it is some fraction of mean or median income) then an increase in the income of a non-poor person (weakly) increases the value of poverty measures. This is a natural property of an inequality measure (Atkinson 1970).

The fact that measures such as the poverty gap either increase or are unchanged when the incomes of the non-poor rise seems natural if the problem is seen as one of measuring inequality with a particular focus on the poor. The behavior expected of a measure of the difficulty of ending poverty is quite different. Obviously, an improvement in the incomes of poor people should reduce the measure, as with standard poverty measures. On the other hand, an increase in the income of the non-poor should also reduce the measure, though not by as much. This is because an increase in the income of the non-poor provides additional resources some of which may be applied to the task of reducing poverty.

3. The poverty burden

In this section, we introduce and characterize a measure of the difficulty of ending poverty. We consider a population of N individuals, with incomes $(y_1 \ldots y_N)$, and a predetermined poverty line z. We assume that the set of individuals with income $y_n > z$ is non-empty. We consider the class of indexes $P(y_1 \ldots y_N, z)$, depending only on the income vector $(y_1 \ldots y_N)$ and the poverty line z.

The Poverty Burden, denoted PB is given by

$$PB\ (y_1 \ldots y_N, z) = (\Sigma_{n \in N} \max(z\text{-}y_n,0))\ /\Sigma_{n \in N} \max\ (y_n\text{-}z,0)$$

The numerator is the poverty gap. The denominator is total income in excess of the poverty line and represents the resources potentially available to reduce poverty.

The Poverty Burden has the following properties:

P.1 Decomposability Given subpopulations $N_i = N_1 \ldots N_I$, the index for the aggregate population is equal to

$$P = \Sigma_i\ s_i P_i$$

where

$$s_i =\ \ (\Sigma_{n \in N_i} \max\ (y_n\text{-}z,0))/(\Sigma_{n \in N} \max\ (y_n\text{-}z,0))$$

is the share of subpopulation N_i in total income in excess of the poverty line and P_i is the index for group N_i.

P.2 Zero in the absence of poverty If $y_n \geq z\ \forall\ n$, $P = 0$.

P.3 Normalized to unity for a poor society   If $\Sigma_{n \in N} y_n / N\ = z$, $P = 1$.

Decomposability is an appealing property for a measure designed for policy implementation. In particular, it ensures that given two populations for which the index is equal to some Poverty Burden, the index for the combined population is also P. More generally, the index for an aggregate population is a weighted average of the indexes for the subpopulations.

Indexes with the property of decomposability, such as those used by Anand and Kakwani have been criticized because they are insensitive to transfers among the poor (and also to transfers among the non-poor). In the present context, this property is an advantage. The Poverty Burden measure is designed with a focus on transfers from the non-poor to the poor, and not on the distribution of income within these groups.

The Zero property is a natural requirement for any measure of poverty. If there are no poor people, then the measure of poverty should be equal to zero.

Finally, P.3 is a natural normalization for a measure of the magnitude of the poverty problem. If the income of a country is so low that, even distributed equally, it cannot lift everyone above the poverty line, it seems reasonable to regard the country as a whole as being poor. Note that the definition of the index does not preclude values above 1, for cases when the average income is below the poverty line.

We now derive our main result:

Proposition 1: *The Poverty Burden **PB** is the unique index satisfying P.1 to P.3*

> **Proof:** We first show that an index P satisfying P.1–P.3 is independent of the distribution of income among the non-poor. Consider any two[1] non-poor individuals j, k. By P.2, $P(y_j, y_k, z) = 0$. Now consider a transfer δ from j to k that leaves both individuals non-poor, so that

$$P(y_j\text{-}\delta,\ y_k + \delta, z) = 0.$$

Let Y_{-jk} denote the set $\{y_1 \dots y_N\}$ with y_j and y_k removed. Set

$$s = ((y_j\text{-}z)+(y_k\text{-}z))/(\Sigma_{n\in N} \max (y_n\text{-}z,\ 0))$$

Then, by P.2,

$$P(y_1 \dots y_\char`^, z) \qquad = s\, P(y_j\text{-}\delta,\ y_k + \delta,\ z)+(1\text{-}s)\, P(y_{\text{-}nm}, z)$$

---

[1] The proof requires that the number of non-poor individuals should be greater than 2. The cases when there are one or two non-poor individuals are trivial.

$$= (1\text{-}s)\, P(y_{\text{-}nm}, z)$$

$$= sP(y_j\text{-}\delta,\ y_k\text{-}\delta, z) + (1\text{-}s)\, P(y_{\text{-}nm}, z)$$

$$= P(y_1 \ldots y_j\text{-}\delta \ldots y_k + \delta \ldots y_N,\ z)$$

where the second equality holds because $P(y_j\text{-}\delta,\ y_k + \delta,\ z) = 0$ since both are non-poor after the transfer. By a similar argument it may be shown that the index is independent of the number of non-poor people.

We now prove the main result for a population such that average income exceeds the poverty line. Let the number of poor people be M. From the preliminary arguments, we may assume, without loss of generality, that, for each poor person with income $z\text{-}\delta$ there exists a non-poor person with income $z+\delta$.

Divide the population into $M+1$ subpopulations, with $1 \ldots M$ consisting of one poor and one non-poor person selected to have total income $2z$ as above, and $M+1$ consisting of all the remaining non-poor persons. Then:

$$P(y_m,\ z) = 1,\ m = 1 \ldots M,$$

$$P(y_{M+1},\ z) = 0$$

and hence

$$P(y_1 \ldots y_N,\ z) \quad = \Sigma_{n \in N}\, s_m,$$

$$= \quad (\Sigma_{n \in N i}\, \max\,(y_n\text{-}z,\ 0)) / (\Sigma_{n \in N}\, \max\,(y_n\text{-}z,\ 0))$$

which is the poverty index PB. ∎

A closely related alternative to the poverty burden measure is based on the tax rate required to be applied to all income rather than income in excess of the poverty line. That is, we seek to determine

$$PB^* = t/(1+t)$$

where t is the rate of tax on all incomes required to fund sufficient transfers to raise all post-tax and transfer income to the poverty line.

To derive this obtain P.1* by replacing s_i in P.1 with

$$s^*_i = (\Sigma_{n \in N_i} y)/(\Sigma_{n \in N} y_n))$$

We obtain

Corollary 1.1: *The unique index satisfying P.1*,P.2, P.3 is*

$$PB^* = 2PB/(1+PB)$$

An increase in income, whether it accrues to the poor or the non-poor, reduces the poverty burden PB. Hence, defining

$$WB(y_1 \ldots y_N, z) = 1 - PB(y_1 \ldots y_N, z)$$

we obtain a social welfare function. the properties of which may be characterized as follows:

Proposition 2: *The index WB is a social welfare function which is*
(i) Strictly increasing in y
(ii) Strictly decreasing in z
(iii) Homothetic of degree zero in its arguments
(iv) Strictly egalitarian

All of these properties apply, with a change of sign, to PB.

   Proof: (i)–(iii) are trivial. For (iv) we observe that if m is poor

$$\partial W/\partial y_m = (\Sigma_{n \in N} \max(z-y_n, 0))+(\Sigma_{n \in N} \max(y_n-z, 0))^2/(\Sigma_{n \in N} \max(y_n-z, 0))^2$$

$$= (1+PB)/(\Sigma_{n \in N} \max(y_n-z, 0)),$$

while if m is non-poor

$$\partial W/\partial y_m = (\Sigma_{n \in N} \max(z-y_n,0))/(\Sigma_{n \in N} \max(y_n-z,0))^2$$

$$= PB/(\Sigma_{n \in N} \max(y_n-z,0)).$$

Thus the welfare function may be locally approximated by a concave angle function. As observed by Machina (1982), this is sufficient to ensure that the function is globally concave and hence egalitarian.∎

Briefly commenting on the properties derived above, monotonicity, derived as 2(i), is a standard feature for a social welfare function. In the context of poverty discussions, it seems appropriate that the higher the poverty line the lower the social welfare associated with any given distribution. Homotheticity 2(iii) is a desirable property for index numbers in general. In the present context, it ensures that changes in the unit of measure do not affect the value of the index. Further, if the concept of poverty is defined in strictly relative terms (for example as some proportion of mean or median income) a proportional increase in all incomes leaves the index unchanged. Finally, 2(iv) (egalitarianism) is a natural requirement.

4.Effectiveness and fungibility

It is inevitable, even if policies are designed to focus on the objective of alleviating poverty, that not all aid will flow to poor people. Administrative costs and incomplete targeting are inevitable, as are, to some degree, rent-seeking and corruption. More importantly, as noted above, many governments adopt aid policies that are not focused on alleviating poverty.

Secondly, aid may have counterproductive effects. These include incentive effects familiar from debates in developed countries over the redistributive effects of tax–welfare systems. In addition, aid delivered in kind, such as food aid, may reduce the incomes of domestic producers who may, in some cases, be poorer than the aid recipients. However, there is no reason to believe that such second-round effects will be such as to more than offset the benefits of aid.

In the literature on the relationship between aid and poverty, summarized by Rajan and Subramanian (2008), discussion of the effectiveness of aid has focused

almost exclusively on the effects of aid on subsequent economic growth. This focus is implicit in the term 'development aid'. This focus contrasts sharply with the discussion of domestic welfare policies, where the primary focus is on raising the consumption levels of households with low (or no) market income.

A focus on aid as a driver of economic growth is problematic in a number of respects. First, as Rajan and Subramanian (2008, p659) observe:

> even in theory—unless one postulates significant nonlinearities in the relationship between investment and growth—it is unlikely that resources for investment by themselves, no matter how well used, will lead to substantial increases in growth[2].

The reason is straightforward. Even assuming that all aid were allocated to investment, the prediction coefficient of aid (expressed as percentage points of GDP) in a growth regression is unlikely to be more than 0.16 and is more plausibly estimated at around 0.03, the typical coefficient on the investment to GDP ratio in cross-section growth regressions derived by Barro and Sala-i-Martin (1995, chapter 12). These coefficients are within the standard error of most of the regressions reported by Rajan and Subramanian. Hence, the statistical tests available to us do not permit us to distinguish between the null hypothesis that aid has no effect on growth and the theoretical prediction derived from the assumption that all aid is invested.

Aid effectiveness and the persistence of poverty

Going further, it is commonly argued that the fact that poverty has persisted in many countries, despite large infusions of aid, shows that aid must be ineffectual (Easterly 2007).

---

[2] Rajan and Subramanian add the qualification ' unless accompanied by policies that increase total factor productivity'. But this qualification is applicable only on the hypothesis the adoption of these policies is induced by the receipt of foreign aid.

Consideration of the poverty burden shows that such arguments cannot be sustained. As previously noted, development aid provided by wealthy countries has never approached the agreed target of 0.7 per cent of income and has generally been around 0.3 per cent of world income. By contrast, until the 1990s, the poverty burden exceeded 3 per cent of world income.

Hence, even on the maximally optimistic assumption that all aid flowed to the world's poorest people, with no adverse second-round effects, the aid that has been provided since the mid-20th century would not have been sufficient to generate a substantial reduction in poverty.

This argument is applicable whether aid is allocated to increasing the consumption of the poor or (less plausibly) to investment projects whose benefits accrue exclusively to the poor. Investment projects as generating a flow of consumption benefits. As the discussion of Rajan and Subramanian implies, for any plausible estimate of rates of return, the volume of aid actually provided to poor countries has been insufficient to generate returns that would produce a substantial reduction in poverty.

But the criterion for the effectiveness of aid implicit in the argument under discussion is even more ambitious than this. The implicit assumption is that if people, having received and consumed aid in the past, remain poor, then the aid has been ineffective. In the context of domestic anti-poverty policy, the analogous test would be that social assistance should not merely have no adverse incentive effects but should induce increases in the market income of the poor so great as to eliminate poverty. On such a criterion, all forms of income assistance, and most other forms of social assistance would certainly fail.

Thus, consideration of poverty burden measures shows that widely used arguments for the ineffectiveness of aid fail basic tests of arithmetic consistency.

Fungibility and the proportion of aid received by the poor

The poverty burden is a lower bound for the cost, relative to aggregate income, of a set of transfers sufficient to raise all households to the poverty line. Although such a lower bound is of conceptual interest, a more relevant issue for policy is the cost of a *feasible* program of transfers. The crucial question here is what proportion of aid is received by the poor. This question in turn depends on the extent to which aid directed to the poor may be treated as fungible income by the government concerned, and to what extent the poor benefit from increases in such fungible income.

To the extent that aid is only partially effective in raising the incomes of poor people the transfer required to achieve any given poverty elimination goal will be increased. That is, if the proportion of aid that is received by the poor is denoted by θ, where $0 < \theta < 1$, the level of aid required to eliminate poverty is PB/θ.

This way of framing the issue suggests a number of possible approaches to estimating θ. The ideal case $\theta = 1$ gives a lower bound to the cost of eliminating poverty, but is of limited practical relevance.

Alternatively we might assume that aid is used to provide public goods, with an equal benefit to all members of the population. In this case θ is equal to the headcount measure of poverty, that is, the proportion of the entire population that is poor.

Discussion of the proportion of aid reaching the poor has been couched mainly in terms of 'fungibility'. The fungibility literature implies that aid might be allocated in proportion to existing public expenditure. The central idea is that aid directed to the poor by donor countries[3] may be used to replace existing

---

[3] In practice, much aid is not directed at improving the living standards of the poor, but at meeting other objectives of the donor country, such as provision of support to friendly governments. However, in considering the possible use of the poverty burden measure by aid agencies in donor countries, it seems appropriate to focus on the question of how much aid,

government expenditure, with the resulting additional revenue being used either to fund expenditure in other areas (Burnside and Dollar 2000)

The allocation of expenditure is generally pro-poor in the sense that poor people receive more, on average, in welfare payments and government services than do those with higher incomes. Thus, on this assumption, the value of θ would be bounded below by the headcount.

A less favorable assumption, adopted by Collier and Dollar (2002) is that the benefits of aid take the form of a proportional increase in income for all members of the community. This would be the case on the joint assumptions that

(i) aid is fully fungible

(ii) additional government income is used entirely to reduce taxation, and

(iii) that the total tax burden is proportional to income

In this case, θ is equal to the proportion of income received by the poor, and the poverty-reducing effects of aid arise only from the fact that income is transferred from rich to poor governments, in a form of fiscal equalization.

Considering the realism of the assumptions required to derive Collier and Dollar's conclusion, (iii) is approximately correct for most taxation systems, but (i) and (ii) seem overly pessimistic.

6. The   magnitude of the poverty burden

In this section, we present estimates of the Poverty Burden for the world and various regions,   using $1/day and $2/day poverty lines. Data in 2005 values for the poverty gap ratios were obtained from the World Bank using the PovcalNet

---

directed at reducing poverty, would be required to raise all poor people to the poverty threshold.

calculations[4] and GDP per capita data were obtained from World Development Indicators.

Table 1:   Poverty Burden (per cent) Using $1/day Poverty Line

	1981	1984	1987	1990	1993	1996	1999	2002	2005
East Asia & Pacific	51.2	29.1	19.0	15.7	11.0	5.5	5.0	3.1	1.3
Europe & Central Asia	n.a.	n.a.	n.a.	0.1	0.3	0.3	0.3	0.2	0.1
Latin America & Caribbean	0.7	0.8	0.8	0.6	0.6	0.6	0.7	0.6	0.5
Middle East & North Africa	0.4	0.3	0.3	0.2	0.2	0.2	0.2	0.1	0.2
South Asia	26.0	21.4	18.6	15.2	12.2	10.5	8.8	8.0	5.9
Sub-Saharan Africa	15.6	18.0	18.3	19.9	20.7	20.5	20.2	17.9	14.6
World	4.3	3.3	2.7	2.5	2.3	1.8	1.7	1.5	1.1
India	26.1	21.2	18.1	14.5	12.8	10.0	8.4	7.5	5.7
China	86.4	42.1	23.1	22.5	14.0	6.4	5.4	3.3	1.2
Bangladesh	22.3	22.3	22.3	23.4	23.2	21.3	22.9	19.5	15.9
World excluding India & China	1.6	1.5	1.5	1.3	1.2	1.2	1.1	1.0	0.8

Source: World Bank, authors' calculations

Table 2:   Poverty Burden (per cent) Using $2/day Poverty Line

	1981	1984	1987	1990	1993	1996	1999	2002	2005
East Asia & Pacific	114.1	65.1	47.9	40.1	29.4	18.2	15.8	10.7	6.0
Europe & Central Asia	n.a.	n.a.	n.a.	0.5	0.9	1.2	1.3	1.0	0.6
Latin America & Caribbean	2.1	2.5	2.2	1.9	1.8	1.9	1.9	1.9	1.3
Middle East & North Africa	3.2	2.5	2.7	2.3	2.2	2.1	1.9	1.7	1.5
South Asia	67.0	59.3	53.7	46.0	40.8	35.7	30.9	28.6	22.5
Sub-Saharan Africa	31.8	35.4	35.7	37.2	39.7	39.6	39.2	36.4	31.2

---

[4] See http://iresearch.worldbank.org/PovcalNet

World	10.1	8.0	7.0	6.5	6.3	5.5	5.1	4.4	3.5
India	67.4	59.5	53.8	45.3	42.0	34.6	29.9	27.2	21.4
China	195.7	93.8	60.3	54.8	36.4	21.0	16.3	10.7	5.4
Bangladesh	84.3	82.4	79.5	78.5	74.2	66.6	64.8	57.6	49.9
World w/o India and China	3.6	3.6	3.4	3.0	3.0	3.0	2.9	2.6	2.2

Source: World Bank, authors' calculations

Tables 1 and 2 have a number of striking features. For the world as a whole, the Poverty Burden using the $1/day poverty line declined between 1981 and 2005, from 4.3 per cent to 1.1 per cent. The decline using the $2/day poverty line was even larger in absolute terms, from 10.1 per cent to 3.5 per cent.

The calculations reported in Table 1 show that, if all countries contributed taxation revenue equal to a little over 1 per cent of national income to global poverty alleviation, and if this revenue flowed to the world's poorest people with no leakages or adverse second-round effects, extreme poverty could be eliminated. An allocation of 3.5 per cent of income, comparable to the Marshall Plan at its peak would be equal to the entire amount required to raise minimum incomes to $2/day.

The decline in the Poverty Burden is primarily the result of favorable developments in East and South Asia. Between 1981 and 2005, the poverty gap in East Asia and the Pacific declined from 51.2 per cent to 1.3 per cent for the $1/day poverty line and 114.1 per cent 6.0 per cent for the $2/day line. The results for the $2/day poverty line are particularly striking In 1981 the region as a whole had an average income below the $2/day poverty line. By contrast, by 2005, the poverty burden was small enough that it could be addressed using the

region's own resources on the basis of expenditures smaller, in relation to national income than those typical of welfare systems in developed countries.

The decline in South Asia is less dramatic, but still notable, for both measures. By contrast, the poverty burden in Sub-Saharan Africa has barely changed. This lack of progress is the result of stagnation in income per person, combined with very little change in overall income distribution.

Finally, the results for the world without India and China are of some interest. Given the rapid economic growth experienced in these two countries, at least until the onset of the global financial crisis in 2008, it seems reasonable to suggest that extreme poverty in these countries could be addressed primarily through domestic policies of redistribution, regional development assistance and so on.

On the $1/day measure, the Poverty Burden for the world without India and China was only 0.8 per cent in 2005, very close to the long-standing target for development aid of 0.7 per cent of national income/GDP, reaffirmed in Monterrey at the 2002 United Nations International Conference on Financing Development.

Although there is a significant gap between the commitment of 0.7 per cent and the poverty burden of 1.1 per cent, much of this gap reflects the persistence of extreme poverty in some rural areas of India and China. If poverty in these rapidly-growing countries were addressed from domestic resources (including continued growth, internal remittances and domestic redistribution), and the developed world adhered to its commitments, the resulting revenue would be sufficient to meet the goal of poverty elimination for the remaining countries where poverty is high (mainly in Sub-Saharan Africa and South Asia). That is, if China and India resolved their own poverty problems, and aid were targeted entirely at poverty reduction, the amount promised at Monterrey would be sufficient to eliminate extreme poverty.

This point can be made more precise, using the Poverty Burden measure. The Poverty Burden for the world, excluding China and India, was 0.78 per cent in

2005, only marginally in excess of the 0.7 per cent target. Assuming a continuation of the declining trend observed in Table 1, it seems reasonable to conclude that, if developed countries met their obligations, and aid were effectively targeted at the poor, extreme poverty could be eliminated within a few years.

Unfortunately, few developed countries have met their commitments, and many have directed aid primarily to governments seen as important for geopolitical reasons, rather than to poor people.

Few nations are close to the agreed 0.7 per cent target for development aid. Among developed countries, only Norway, Denmark, Luxembourg, The Netherlands and Sweden exceeded the target. Belgium and Austria were the only other countries to provide even 0.5 percent of GNP. The average overseas development aid to donor GNI was 0.25 per cent in 2003 (United Nations Millennium Project 2005).[5]

Concluding comments

Measures of poverty and inequality are useful diagnostic tools in assessing the equity, or otherwise, of income distribution. In this paper, it has been argued that a different class of measures, referred to as poverty burden measures, is needed to assess the magnitude   of the task of eliminating extreme poverty.

Taking considerations of effectiveness and targeting into account, it seems likely that the volume of additional income required to end extreme poverty would be significantly in excess of the value of 1.1 per cent of world income estimated in Table 1. Nevertheless, if the developed world could agree on transfers comparable to those made by the United States during the Marshall Plan era (around 2 per cent of GDP annually), and if global economic growth recovers quickly from the financial crisis, a rapid reduction in extreme poverty should be achievable.

---

[5] The ratio of overseas development aid to GDP for the developed countries was 0.51 per cent in 1960, 0.33 per cent in 1970, 0.35   per cent in 1980, 0.34   per cent in 1990 and 0.23   per cent in 2002 (UN 2005).

Aid must be allocated between consumption (aid that provides direct improvements in current living standards) and investment (aid that is aimed at increasing living standards in the future). Much policy discussion implies that investment is always to be preferred to consumption. But economic analysis gives no basis for such a preference. Increasing the current consumption of the poor is good in itself, whether or not it contributes to increases in future income and consumption.

References

Anand, S. (1977): Aspects of Poverty in Malaysia, Review of Income and Wealth, 23, 1–16.

Atkinson, A.B. (1970): On the Measurement of Inequality, Journal of Economic Theory, 2, 244–63.

Beckerman, W. (1979): The Impact of Income Maintenance Payments on Poverty in Britain, 1975, The Economic Journal, 89, 261–79.

Blackorby, C. and Donaldson, D. (1980): Ethical Indices for the Measurement of Poverty, Econometrica, 48, 1053–60.

Burnside, C. and Dollar, D. (2000): Aid, Policies, and Growth, American Economic Review, 90, 847-68.

Collier, P. and Dollar, D. (2002): Aid Allocation and Poverty Reduction, European Economic Review, 46, 1475–500.

Dalgaard, C., Hansen, H. and Tarp, F. (2004): On the Empirics of Foreign Aid and Growth, Economic Journal, 114, F191–F216.

Easterly, W. (2007): The White Man's Burden: Why the West's Efforts to Aid the Rest Have Done So Much Ill and So Little Good, Penguin, London.

Foster, J., Greer, J. and Thorbecke, E. (1984): A Class of Decomposable Poverty Measures, Econometrica, 52, 761–66.

Foster, J.E. (1998): Absolute Versus Relative Poverty, American Economic Review, 88, 335–41.

Foster, J.E. and Shorrocks, A.F. (1988): Poverty Orderings, Econometrica, 56, 173–77.

Foster, J.E. and Shorrocks, A.F. (1991): Subgroup Consistent Poverty Indices, Econometrica, 59, 687–709.

Kakwani, N. (1980): On a Class of Poverty Measures, Econometrica, 48, 437–46.

Machina, M. (1982): 'Expected Utility' Analysis Without the Independence Axiom, Econometrica, 50, 277–323.

Rajan, R. and Subramanian, A. (2008): Aid and Growth: What Does the Cross-Country Evidence Really Show?, Review of Economics and Statistics, 90, 643–65.

Ravallion, M. (1996): Issues in Measuring and Modelling Poverty, Economic Journal, 106, 1328–43.

Sachs, J. (2005): The End of Poverty: Economic Possibilities for Our Time, Penguin, New York.

Sen, A. (1976): Poverty: An Ordinal Approach to Measurement, Econometrica, 44, 219–31.

Shorrocks, A.F. (1995): Revisiting the Sen Poverty Index, Econometrica, 63, 1225–30.

Shorrocks, A.F. and Foster, J.E. (1987): Transfer Sensitive Inequality Measures, Review of Economic Studies, 54, 485–97.

United Nations Millennium Project, (2006): The 0.7% Target: An in-Depth Look, United Nations, New York.

Vaughan, R.N. (1987): Welfare Approaches to the Measurement of Poverty, Economic Journal, 97, 160–70.

World Bank (2008): World Development Report, 2008, World Bank, Washington, DC.

Zheng, B. (1994): Can a Poverty Index be Both Relative and Absolute?, Econometrica, 62, 1453–58.