



Initiative for **Policy Dialogue**

COPYRIGHT NOTICE: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior permission in writing of the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to share this file through any network server.

1

Introduction

Sudhir Anand, Paul Segal, and Joseph E. Stiglitz

Global poverty is higher on the public agenda than ever before. While it has always been a central concern of development economists and professionals, it has never before been studied so intensively, and has never been as prominent in the public consciousness as it has become over the last few years. This new focus is due in part to the commitment of the international community to the United Nations Millennium Development Goals (MDGs)—the first of which is to halve, between 1990 and 2015, the proportion of people living in extreme poverty. It is due also to substantial progress in the collection and analysis of data. While we can now say much more about global poverty than in the past, considerable controversy has developed about how to analyze and interpret the data, and what they tell us about the magnitude and rate of change of global poverty.

In many ways, the controversy should not come as a surprise. If the world's attention were not focused on poverty, we would have little need to measure it. Differences among economists and statisticians in the estimates would be of little significance. But with a global commitment to halve poverty, how we measure it makes a great deal of difference. If in one system of metrics we have already achieved our goal, the MDGs will not be able to serve as the rallying cry for more assistance (whether of aid or trade). At least some of the moral weight behind the demands by civil society is that the persistence of poverty at this level is unconscionable in a world of such wealth.

The issue of measurement is thus of more than just academic interest. What we measure affects what we do. The focus on poverty has directed scarce resources towards poverty reduction. Many of those who believe that we have made great strides in reducing poverty—that we have already achieved the MDG of halving poverty since 1990—believe that these scarce resources should be devoted to maximizing growth. Other critics of poverty measurement worry that focusing excessive attention on numbers in poverty encourages governments to direct resources towards those just below the poverty line as this is the easiest way to reduce the number of poor. But this is not the best way to improve the well-being of the poor.

Introduction

All of the poverty lines used to measure global poverty represent living standards that are hard to fathom by those in advanced industrial countries. What would it be like to live on \$1 a day, or \$2 a day—on an annual income of \$365 to \$730? For most of us, it is beyond conception. The poverty standards in the US are closer to \$15 a day, and a visit to the slums of Detroit or New York provides a picture of what life is like for those with incomes at this low level. Having applied for a graduate fellowship very late in his third year in college, and without a degree, Stiglitz was fortunate enough to receive a small fellowship from MIT which left him with just \$1 a day for food (in 1963 prices—equivalent to \$7 in today's prices). Even with careful planning, it was not sustainable: Stiglitz lost weight rapidly during the year (which he was happily able to regain subsequently). For Stiglitz it was an experiment—one which he knew would end in twelve months. But for those in the developing world, it is not an experiment, and for many no end to starvation level diets is in sight. There can be a vicious circle with poor diets leading to low productivity, and low productivity leading to low wages (see, for example, Stiglitz, 1976).

Every once in a while, the neglected field of national income statistics finds itself at the center of controversy. For example, the measurement of inflation in Argentina has been disputed in recent years: has the government cooked the books so that true inflation is higher than measured inflation? Is it trying to allay the longstanding concerns of that society about inflation, not by lowering inflation, but by lowering the measurement of inflation? Ten years ago, when Stiglitz served as Chair of the Council of Economic Advisers in the US, another such controversy flared up. Were we overestimating inflation, thereby giving social security recipients (whose payments go up with inflation) far too much year after year, and undermining the financial viability of the entire social security system? A report by a former Chair of the Council of Economic Advisers (Boskin Commission, 1996) found that inflation was overestimated by 1.3 percentage points per year from 1978 to 1996.

In some cases, there is a strong rationale for measuring things one way or another. In others, it is simply a matter of convention—though if we are to make comparisons across countries or over time, the comparisons must be done in a consistent way. With prices and products changing all the time and differing across and within countries, making consistent comparisons is harder than it might seem, and some of the controversies relate to these differences. With brie and croissants relatively more expensive in the US, and McDonalds hamburgers relatively more expensive in France, a Frenchman might find his cost of living in the US far higher than in France; but at the same time, an American might find maintaining his life style in Paris more difficult than in the US. Economists have, over the years, developed standard, if imperfect, techniques for dealing with these problems.

Even if there were not these price and product differences, there is another question. How do we define poverty? Some of the original poverty measures

were devised as basic measures of survivability. Below those levels, it was hard to sustain life—to meet nutritional needs and provide the basic necessities of shelter, energy for cooking and heating, and clothing. Of course, these basic necessities could differ in different parts of the world—illustrating that a global standard (e.g., \$1 or \$2 a day) must be viewed simply as an approximation, a point of reference. It should be noted, though, that measured poverty corresponding to different poverty lines can show contrary trends depending on changes in the income distribution. Moreover, updating the poverty line over time in a world of changing products and prices leads to further challenges.

Beyond national income statistics, there are also serious issues regarding the household surveys that are used to measure living standards within countries. Most data on the distribution of income come from asking people their income and consumption levels. It is often difficult for individuals to recall this information perfectly. We know that something is wrong because the total amount reported, when extrapolated to the whole economy, does not add up to the national accounts estimate of total income of the economy. Is this due to conceptual differences between surveys and national accounts, or to underreporting in surveys? If the latter, is the underreporting uniform, or are the rich underreporting more than the poor? The answer to these questions, as we shall see, makes a great deal of difference to one's view about what is happening to global poverty.

This volume brings together some of the leading researchers in the field, from both academia and international organizations, in order to provide a thorough examination of the challenges and uncertainties involved in measuring global poverty. The chapters in Part I of the book discuss questions on the measurement of poverty at a global level, including conceptual issues such as what we really mean by poverty—and whether we can apply the concept at the global level—to the more empirical questions of whether to use national accounts data or household survey data, and the difficulties posed by the use of purchasing power parity (PPP) exchange rates. Part II then presents studies on specific regions and countries of the world in order to provide an overview of the challenges faced.

Estimates of global poverty

The standard estimates, most commonly cited by academics and the media, are due to Shaohua Chen and Martin Ravallion (2001, 2004, 2007, 2008) at the World Bank. Using the World Bank's "PPP\$1-a-day" poverty line,¹ Chen and Ravallion's (2008) estimates are that 1,377 million people were living in extreme poverty in 2005, compared with 1,813 million in 1990 (see Table 1.1).² Poverty is going down, but slowly. At this pace, assuming a

Introduction

Table 1.1. Estimates of global poverty, 1980–2005 (millions)

Year	1980	1990	2000	2005
Chen and Ravallion (2008) PPP\$1.25, 2005 prices	1,896 ^a	1,813	1,696 ^b	1,377
Chen and Ravallion (2007) PPP\$1.08, 1993 prices	1,470 ^a	1,248	1,109 ^b	969 ^d
Bhalla (this volume, Ch. 4) “SNAK” PPP\$1.08, 1993 prices	1,489 ^a	1,216	770 ^c	456
Bhalla (2002) adjusted HFCE PPP\$1.50, 1993 prices	1,479	1,056	647	
Bhalla (2002) survey means PPP\$1.30, 1993 prices	1,581	1,208	899	
Sala-i-Martin (2006) PPP\$1.36, 1996 prices	498	363	322	

^a 1981 estimate;

^b 1999 estimate;

^c 2001 estimate;

^d 2004 estimate.

linear trend in global poverty reduction would yield a total of 1,086 million poor in 2015. But performance varies widely across regions. China has had enormous success in reducing poverty. In Africa, however, the prospects of meeting the MDGs are bleak, with poverty going up from 299 million in 1990 to 391 million in 2005.

The World Bank’s canonical estimates have been challenged by researchers from both the left and the right. Some believe that the World Bank numbers underestimate the level of poverty and overestimate its reduction (see, in particular, Reddy and Pogge, Ch. 3 in this volume); while others (Surjit Bhalla, Ch. 4 in this volume, and Sala-i-Martin, 2006) believe that the World Bank numbers overestimate the level of poverty and underestimate its reduction. Table 1.1 and Figure 1.1 present estimates of global poverty by Chen and Ravallion (2007, 2008), by Bhalla (2002, and this volume), and by Sala-i-Martin (2006). As can be seen, their estimates vary substantially. Chen and Ravallion’s 2007 estimates are much lower than their 2008 estimates, but their rates of poverty decline are not very different. Bhalla in his chapter, ostensibly using the same World Bank poverty line, estimates the number in extreme poverty to decline dramatically from 1,216 million in 1990 to only 456 million in 2005. Sala-i-Martin (2006) estimates global poverty at just 363 million in 1990 and 322 million in 2000—a slower rate of decline, but from a much lower level. In the context of trying to achieve the first MDG, our views on both the extent of global poverty and its rate of decline have great policy relevance. Is a concerted international effort required? Or is it true, as Sala-i-Martin claims, that “[t]he world might just be in better shape than many of our leaders believe!”³




Figure 1.1 Estimates of global poverty, 1980–2005 (millions)

The measurement of global poverty

The large divergence in estimates of global poverty arises from the use of different methodologies and datasets, which we now describe in greater detail. First, the authors use different poverty lines. Second, they use different PPP exchange rates to convert incomes in local currencies into a common international currency.⁴ Third, they take different approaches to estimating within-country distributions of income. Fourth, they calculate mean incomes within countries differently.

The definition of poverty

The first question in the measurement of global poverty is: what do we mean by poverty? Put another way, we have to define a poverty threshold. The income poverty lines used by Bhalla, Sala-i-Martin, and the World Bank (i.e. Chen and Ravallion) are all derived, in some sense, from the original World Bank PPP\$1-a-day poverty line at 1985 prices (World Bank, 1990). This poverty line was informally chosen as being representative of the subsistence poverty lines of the poorest countries, translated into 1985 PPP\$. Yet, while all authors base their poverty line on the same concept, they adopt different numbers. Chen and Ravallion (2007) use PPP\$1.08 at 1993 prices.

Introduction

Bhalla (2002) states that the “equivalent to \$1 a day at 1985 prices is \$1.30 a day at 1993 prices and not \$1.08 a day” (p. 67). He thus uses a line of \$1.30 at 1993 prices with survey data, while his “most preferred” method (p. 140) is to use national accounts data and adjust the poverty line to \$1.50.⁵ In his chapter in this volume (Ch. 4), Bhalla uses the Chen and Ravallion (2007) line of \$1.08 at 1993 prices. Sala-i-Martin (2006) states that the poverty line of \$1 a day at 1985 prices is equivalent to PPP\$1.36 per day, or PPP\$495 per year, at 1996 prices. In their most recent update based on a new global price survey (discussed below), Chen and Ravallion (2008) adopt a line of PPP\$1.25 at 2005 prices. How can these differences be explained?

The disagreement arises over how to update a 1985-based PPP\$ value. Within a single country, one would usually update a poverty line using a price index based on measured inflation. Thus if prices are estimated to have risen by 10 per cent after the poverty line was established, then the poverty line today should be 10 per cent higher. If we are expressing everything in US dollars, then all we need to do is to look at US inflation over the period 1985 to 1993. But updating a poverty line that is denominated in PPP\$ is not so simple. Calculating a set of PPP exchange rates involves the prices of *all* countries (discussed below), so changes in a country’s PPP exchange rate will depend on price changes in all countries. Bangladesh’s 1985 poverty line at 1985 PPP\$ scaled up by US inflation over 1985–93 would not be expected to be equal to Bangladesh’s 1993 poverty line at 1993 PPP\$.⁶

This is part of the more general “index number problem” that there is no simple, or uniquely best, way to convert overall price levels across space or time when relative prices are changing—and relative prices do change, across both space and time. PPPs are multilateral price indexes, and they have no analytical relationship with the price indexes used within countries to measure inflation. In fact, Deaton (2001, p. 127) observes that “the PPP international dollar has strengthened relative to the currencies of the poor countries whose poverty lines are incorporated into the international line.”

This is indeed what Chen and Ravallion (2001) find. They calculate afresh a global poverty line for 1993 using a similar method to that for the original 1985 poverty line, deriving it as the median of the lowest ten poverty lines in their data set⁷ converted into PPP\$—this time using 1993 PPPs. This re-doing of the original method results in a poverty line of PPP\$1.08 per day of consumption at 1993 PPP dollars, referred to as “\$1 a day” for convenience. This represents a much lower rate of inflation than experienced in the US over 1985–93. As Chen and Ravallion (2001, p. 288) point out, “the fact that \$1.08 in 1993 has a US purchasing power less than \$1 in 1985 does not mean that the real value of the poverty line has fallen. Indeed, if we had simply adjusted the \$1 per day line for inflation in the US between 1985 and 1993 we would

have obtained a poverty line which is well above the median of the ten lowest poverty lines at 1993 PPP.”

Sala-i-Martin (2006) appears to be unaware of this point, however, going so far as to claim that “this mysterious change in the poverty threshold has never been explained by the World Bank” (p. 370). Instead, he bases his calculation of a poverty line on US inflation, stating that the 1985 poverty line corresponds to \$495 a year in 1996 prices (pp. 370, 372). Bhalla (2002) maintains that “international inflation . . . is what is needed to convert incomes (or consumption) from one base to another” (p. 65). Unfortunately, this ignores Deaton’s point regarding the depreciation of the currencies of poor countries relative to international PPP\$. For these reasons the poverty lines used by Bhalla (2002) and Sala-i-Martin (2006) are higher in real terms than those used by the World Bank.⁸

All three of these methods of constructing a global poverty line use a money metric threshold of poverty, converting estimates of household consumption, measured in national currency units, into PPP\$. Reddy and Pogge in this volume challenge the money metric approach to global poverty measurement. Their primary objection to the PPP\$1-a-day poverty line is that it does not correspond to any “achievement concept” or set of capabilities that are common across countries. That is, there is no reason to think that PPP\$1 a day in one country will enable the same set of achievements—e.g. in terms of nutrition or shelter—as PPP\$1 a day in another country. While domestic poverty lines are often set according to some achievement concept, this interpretation is lost when a global poverty line is constructed using standard PPP exchange rates. Reddy and Pogge argue that an explicit achievement-based threshold should be used to define a global poverty line. This would require costing the minimal standard set of capabilities in each country to yield a money-metric poverty line denominated in local currency. Thus the global capability-based poverty threshold would be represented in income space by the set of these national poverty lines, one for each country.⁹

T. N. Srinivasan (Ch. 5) is equally critical of the World Bank’s definition of a global poverty line. He argues that poverty should be seen as a multidimensional concept, but that if it is to be considered in monetary terms it should be defined at the national level only—as the cost of a bundle of basic goods and services specific to each country. He is more pessimistic than Reddy and Pogge regarding the possibility of defining a global poverty line for international comparison and aggregation of poverty across countries. In his view, measures of global poverty will in practice never be satisfactory and are at best just advocacy tools for focusing public attention.

Robert Johnston’s chapter contributes a historical discussion on the measurement of poverty, tracing it back over more than a century. He argues that the concept of extreme poverty as “distress and degradation” developed by early researchers remains appropriate for the measurement of global poverty today.

Introduction

David E. Sahn and Stephen D. Younger's chapter on Africa in this volume follows the achievements-based approach of the human development literature by focusing on non-monetary measures of poverty. As indicators of well-being they use anthropometric measures (height-for-age and body mass index) and years of schooling. The authors make the point that these non-income indicators have a number of practical advantages over monetary indicators, including their being easier to measure reliably,¹⁰ and their being defined at the individual rather than the household level. To these one could add Reddy and Pogge's point that they do not suffer from the difficulties of the PPP\$1-a-day line, and are thus more comparable across countries. On the other hand, they are measures of a longer-term conception of well-being than those based on current real income or consumption, because they depend on nutrition and health, or schooling, over a number of years. Thus they are unlikely to respond to short-run changes in policy or economic environment. Sahn and Younger's data on Africa indicate a mixed picture: no clear trend in health indicators, but a widespread (though not ubiquitous) decline in education poverty.

Sakiko Fukuda-Parr and David Stewart in their chapter similarly provide a multidimensional picture of poverty and development, presenting a broad overview of recent trends in human poverty and human development. Using such wide-ranging indicators as life expectancy and rates of hunger and schooling, in addition to income, they illustrate the point that a full picture requires the use of numerous measures of well-being.

While the above authors use non-monetary measures of deprivation, Riskin and Gao in their chapter on China consider several income poverty lines that are based on the cost of achieving minimum levels of calorie intake. They find that the direction of change of urban poverty between 1988 and 1995 depends on the poverty line used, but there appears to have been a clear decline from 1995 to 2002. The social safety nets put in place by the government do not seem to be behind this reduction, but the government's strategy to make growth more equitable, including its targeted investment programs, may have played a role. Riskin and Gao also analyze a range of factors associated with urban poverty and find, for instance, that while employment in a state-owned enterprise reduced the probability of being poor in 1988, it increased the probability in 1995 and had no impact in 2002. While being unemployed contributes substantially to the probability of being poor, most heads of poor urban households are in fact employed.

Purchasing power parity exchange rates

Assuming one adopts the money-metric approach to poverty measurement, household consumption in national currencies has to be converted into a common currency for international comparison. The use of PPP exchange rates is intended to take account of the fact that, for instance, a dollar's worth

of rupees, bought on the currency markets, will buy more of most goods and services in India than the same dollar would buy in the US. But \$1 converted into rupees at the PPP exchange rate should buy approximately the same quantity of goods and services in India as \$1 does in the US. Incomes in developing countries can be three or four times higher when measured at PPP exchange rates than when measured at market exchange rates.

PPP exchange rates are calculated using price surveys across countries conducted by the International Comparison Program (ICP). But these surveys are not done every year, so the country's real growth rate (measured using its own market basket of goods to adjust for inflation) is used to estimate PPP income in other years. However, real PPP income estimates calculated in this way can turn out to be very different from estimates based on a new price survey (for instance, see World Bank, 2008b).¹¹ Thus estimates of real income using PPPs based on different benchmark years will not be comparable.

ICP price surveys were benchmarked in 1985, 1993–96 (referred to as the 1993 ICP), and 2005. Apart from the latest World Bank estimates (Chen and Ravallion, 2008) which are based on the 2005 ICP, all estimates of global poverty in Table 1.1 use PPPs based on the 1993 ICP. Other researchers have not yet used the 2005 ICP data to generate estimates of global poverty. One advantage of the new ICP round is that a much more comprehensive price survey has been undertaken for China—its PPP exchange rate had previously been estimated on the basis of limited data from a few cities. It is also the first time since 1985 that an ICP price survey has been conducted for India. For both countries, the estimates of GDP using the new PPP exchange rates are substantially lower than previous estimates (World Bank, 2008b). Chen and Ravallion's chapter in this volume discusses the implications of the new price data for China and presents new poverty estimates. China is found to have more poverty than previously thought, but its record in poverty reduction remains just as impressive. This is not, of course, surprising: the new PPP data indicate that income levels in China are lower than implied by previous PPP estimates, but estimates of growth rates are unchanged.

The most widely used set of PPP data are from the Penn World Tables (PWT), and Bettina Aten and Alan Heston in their chapter discuss the use of PWT data in the measurement of global poverty. The PWT PPP estimates are based on the Geary-Khamis (GK) method, in which a vector of "average international prices" is constructed to value the output of each country. Both Bhalla and Sala-i-Martin make use of PWT in their estimations of global poverty.¹² On the other hand, the World Bank, for both its GDP and global poverty estimates, uses the Eltetö-Köves-Szulc (EKS) method. This method does not involve constructing a vector of international prices but is based on bilateral (Fisher) price indices computed using each country's output vector as weights; the PPP exchange rate for a given country is then defined as the geometric average of its bilateral price indices with respect to all other countries.¹³

Introduction

The use by the GK method of a vector of “average international prices” leads to a potential problem. As Aten and Heston (Ch. 6) explain, the “international price” of a good is a weighted average of its price in each country in international dollars, where the weights are the country’s share in world output. This implies that prices in larger economies get a larger weight, so the resulting relative international prices are closer to relative prices in larger economies, which tend to have higher incomes per capita. When incomes are measured at these international prices, substitution bias (the Gershenkron effect)—the fact that people buy more of the goods that are relatively cheap in their own country—implies that the incomes of poor countries are likely to be overestimated relative to the incomes of rich countries. Since the global poverty line is the median of the lowest ten poverty lines at EKS PPP\$ (Chen and Ravallion, 2001), calculating global poverty relative to this line through use of incomes converted at GK PPP will lead to an underestimate.¹⁴ Using the same basic methodology as Sala-i-Martin, a study by Ackland et al. (2004) finds that estimates based on incomes at EKS PPP imply a global poverty incidence nearly 60 per cent higher than those based on incomes at GK PPP (from PWT). In published global poverty estimates, Sala-i-Martin (2006) is alone in making exclusive use of PWT GDP data (based on GK PPPs),¹⁵ and this is likely to be part of the explanation for his much lower estimates of global poverty compared with the World Bank’s EKS-based estimates.

Aten and Heston (Ch. 6) discuss the possibility of “poverty PPPs,” or PPP exchange rates designed specifically to convert the incomes of the poor. They observe that existing consumption PPPs, based on prices only of consumption goods and services, are an improvement on GDP PPPs, which include investment goods and government expenditures. But, as they point out, the consumption basket of the poor will typically be different from the average consumption basket, with food expenditure comprising a larger fraction of the budget of the poor than of the average consumer. The poor may also face different prices from the average consumer. In principle, one could restrict PPP calculations to goods consumed by the poor, and also use prices faced by the poor, but such data are typically not available.¹⁶ These questions regarding the appropriate expenditure weights and prices in the construction of poverty PPPs apply to both the GK and EKS methods.

The fact that standard PPPs are not designed for converting the incomes of the poor is also highlighted by Reddy and Pogge (Ch. 3). They observe that use of these PPPs implies that estimates of global poverty depend on some prices that are irrelevant to the poor. They comment that “whether a household in India lives in absolute poverty by the \$1 PPP per day standard cannot reasonably depend on information about Japanese real estate prices, but under the current methodology of poverty assessment it may.”¹⁷ The use of poverty PPPs would mitigate this problem.

The distribution of income within countries

For measuring global poverty we need an estimate of the distribution of income within each country. Household income or consumption surveys in countries are the main source for this information. World Bank researchers have access to the primary unit record data in such surveys, which enables them to compute a country's income or consumption distribution directly (for convenience, henceforth referred to simply as "income distribution").¹⁸ Other authors who did not have access to unit record data of household surveys have used secondary published information on within-country inequality.

Bhalla (2002) thus uses a two-step method to estimate global poverty. First, he takes estimates of relative inequality within countries from secondary datasets (Deiningen and Squire, 1996, and UNU-WIDER's World Income Inequality Database¹⁹), which provide quintile shares for most countries and decile shares for some. For each country Bhalla scales these income shares to an exogenous estimate of average income or consumption.²⁰ His favored method (2002) is to scale to a constant fraction of the national accounts category of per capita Household Final Consumption Expenditure (HFCE),²¹ but he also provides estimates where income shares are scaled to the mean income as measured in surveys. In his chapter in this volume (Ch. 4), Bhalla presents three sets of estimates. Two estimates are based on survey means from different sources (World Bank and non-World Bank website data respectively). The third method, which he favors, uses the survey mean in 1987 and projects this mean backwards and forwards using the growth rate of consumption from the national accounts, thus keeping constant the ratio of the mean of the distribution to HFCE at its 1987 value; he refers to this method as "SNAk" (estimates reported in Table 1.1 and Figure 1.1). Sala-i-Martin also follows a two-step method, taking quintile share data from Deiningen and Squire (1996 updated dataset) and scaling them to per capita GDP. The quintile numbers are then converted to smoothed within-country distributions using a technique called kernel density estimation.²²

Perhaps the most important methodological difference between alternative estimates of global poverty is the choice of mean for within-country distributions: the survey mean itself, (adjusted) per capita HFCE, or per capita GDP. This choice would seem to account for much of the difference in estimates of levels of and changes in global poverty.

Both Bhalla and Sala-i-Martin scale within-country distributions to national accounts categories, viz. adjusted per capita HFCE and per capita GDP, respectively. One reason given by Sala-i-Martin for scaling to per capita GDP is that he wishes to measure income and not consumption poverty. This, however, would require the use of a national accounts category of aggregate personal or household income. But following the 1993 System of National Accounts, most countries do not include such a category in their published national accounts.

Introduction

GDP is therefore the only national accounts measure of income available across a wide range of countries. However, its use in the measurement of poverty is problematic. GDP includes retained earnings of corporations, the part of government revenue (taxes) that is not distributed back to households as cash transfers, and GDP does not net out depreciation. For illustration we can take the example of the US, which is one of the few countries that does report measures of household income (referred to as personal income) in its National Income and Product Accounts (NIPA). In 2006, US GDP was \$13,246.6 billion, personal income \$10,891.2 billion, and disposable personal income \$9,529.1 billion (NIPA Tables 1.1.5 and 2.1). Disposable personal income was therefore only 72 per cent of GDP. Hence scaling to GDP per capita can substantially overestimate personal incomes and underestimate income poverty.

The Millennium Development Goals monitor consumption poverty,²³ which is what the World Bank estimates. Ravallion (2004) argues that the difference between consumption and income poverty (and Sala-i-Martin's use of GDP to measure income) may account for the difference between Sala-i-Martin's and the World Bank's 1993 PPP\$-based estimates (reported in Chen and Ravallion, 2007).

The category of HFCE does not include those components of GDP mentioned above that are not part of aggregate household or personal income. Moreover, HFCE is reported for all countries in the IMF's *International Financial Statistics*. Hence we find Bhalla's (2002) (adjusted) HFCE-based global poverty estimates are higher than Sala-i-Martin's estimates and closer to the World Bank's (see Figure 1.1 and Table 1.1). However, Bhalla's estimates (SNAk) in this volume based on the growth rate of HFCE decline more rapidly than his or the World Bank's (1993 PPP\$) survey mean-based estimates, and by 2001 they have fallen to approximately half of the World Bank's estimates.

The divergence between survey mean-based and HFCE-based estimates of global poverty is due to a growing divergence between survey means and per capita HFCE. Deaton, Ch. 7 in this volume, finds that, for his sample of non-Organisation for Economic Co-operation and Development (OECD) countries during 1990–2000, “the growth rate of survey consumption is approximately half of the growth rate of national accounts consumption.” In India, for example, the ratio of survey to NA (national accounts) consumption declined over time from 0.68 in 1983 to 0.56 in 1999/2000. There has been heated debate about the source of this divergence in the context of poverty measurement in India (e.g. Ravallion, 2000; Bhalla, 2002; Deaton and Kozel, 2005).

Although both surveys and the national accounts measure “household consumption”, household expenditure in surveys differs from HFCE in the national accounts in both concept and method of estimation. This issue is discussed in detail by Ivo Havinga, Gisèle Kamanou, and Viet Vu in their chapter (Ch. 9). In terms of concept, HFCE includes imputed values of financial intermediation services and consumption by “non-profit organizations serving households.”²⁴ HFCE also includes imputed rents from owner-occupied housing, which are

often not estimated in household surveys. It should be noted that neither survey expenditure nor HFCE includes imputed values of government-provided healthcare, education, or other services.²⁵

The two categories also differ greatly in their method of estimation.²⁶ To calculate HFCE, the national accounts typically starts with an estimate of national production of a commodity such as rice from crop-cutting data, aerial or farm surveys, etc. As such surveys are conducted infrequently, gross production figures may have to be estimated without up-to-date information. Moreover, the methods used to arrive at these figures are not applied uniformly and can be unreliable. From an estimate of national production thus generated, government consumption and firms' consumption are subtracted, and the residual is attributed to households. Data on government consumption may be adequate, but firms' consumption is typically poorly estimated. It is often based on outdated firm surveys, and extrapolations or assumed changes over time. In India, the divergence between survey and national accounts consumption is partly due to the underestimation by NA of firms' consumption of intermediate goods. This has led to double-counting where, for instance, the edible oil consumed in restaurant meals was attributed to HFCE under both the "edible oil" category and the "restaurant meals" category.²⁷

NA estimates of HFCE are thus indirect and subject to three sources of error: the initial estimate of aggregate production, the estimate of government consumption, and the estimate of firms' consumption. There is no reason to suppose that the data and methods used to estimate these, which include surveys of various types, are more reliable than household surveys. Moreover, their sources and methods are generally less well-documented (in terms of the surveys employed, how and when they were conducted, etc.) than household surveys. Finally, because it is calculated as a residual, the errors in the estimate of HFCE will tend to get compounded. By contrast, household surveys measure personal consumption and income directly.

Household income and consumption measured through surveys will, however, also be subject to errors. One reason is that there may be underreporting of incomes. If all incomes were underreported by the same fraction, then uniform scaling to an accurate estimate of mean income would correct the problem. However, the underreporting is not uniform across the distribution. As Ravallion (Ch. 2, this volume) discusses, the rich tend to underreport proportionately more than the poor.

Differential underreporting by the rich and poor can help to explain the growing divergence between mean income or consumption measured in surveys and in the national accounts. As the income share of the rich rises, which has occurred for instance in both India and China (and in other countries where inequality has risen), greater underreporting by the rich compared with the poor will imply growing underestimation of average household income and consumption as measured in surveys.

Introduction

Underreporting by the rich will not by itself imply any bias in the measurement of poverty when individual income or consumption levels are obtained directly from surveys. In this case, individuals above the poverty line are of limited interest, and scaling to national accounts categories can lead to underestimation of the incidence of poverty. In his chapter, Ravallion (p. 33, this volume) provides a simple illustration of this point as follows:

The true but unobserved distribution of income is (say) 1, 2, 3 (person 1 has an income of 1, person 2 has income 2, person 3 has 3). The poverty line is slightly above 1, so the true poverty rate is $1/3$. We do a survey, and the three people respond that their incomes are 1, 1.5 and 2. Income of person 2 is underestimated by one-quarter, while it is underestimated by one-third for person 3. The survey gives the right poverty rate. However, the survey underestimates the true mean; the survey mean is 1.5. Now let's assume (for the sake of argument) that the national accounts do give the right mean of 2. If we assume that the survey under-estimation is distribution-neutral then we multiply all three incomes by $4/3$. The 'corrected' incomes are 1.3, 2 and 2.7—implying that there is no poverty. We get the mean right, but the poverty measure is way off the mark.

Moreover, if a rise in within-country inequality leads to a growing underestimation of the mean, then the extent of underestimation of poverty due to scaling to the NA mean may be increasing over time—which will imply an overestimation of poverty reduction.

A further problem with household surveys is that the respondents may not be representative of the population. The very rich are reluctant to respond for tax and other reasons. Korinek et al. (2006) find that in the US the rich tend to respond less than the poor. In developing countries, on the other hand, the very poor and marginalized—particularly the homeless who have no fixed address, or those living in remote rural areas—may be excluded from the sample frame and are thus also likely to be under-represented.

The implication of under-representation of both the rich and the poor in surveys is different from that of underreporting. In this context even data taken directly from the survey may underestimate poverty. Extending Ravallion's example, suppose that the true distribution is 1, 1, 2, 3, 3, and hence the true poverty rate is $2/5$. Now suppose that half of the poorest are missed from the sample frame and half of the richest fail to respond, so the survey reports only the three incomes 1, 2, 3. The survey mean remains correct at 2, but the survey underestimates poverty at $1/3$ when the true incidence is $2/5$.

Even though surveys are prone to measurement error, there is little reason to think that scaling the mean to some NA category, while using within-country relative distributions from surveys, will reduce the error in measuring poverty. On the other hand, it is doubtful that the decline in survey means relative to NA means could be attributed entirely to underreporting by the rich. Bhalla's method of using the survey mean in 1987, and projecting this mean forwards

and backwards with the growth rate of NA consumption, provides a useful alternative for comparison with the purely survey-based estimates of the World Bank. As seen above, his estimates indicate a more rapid rate of poverty reduction after 1987 than the World Bank's, owing to the growing divergence between survey and NA means.

Clearly there is scope to improve data collection in household surveys. Albert Berry (Ch. 11 in this volume) on the measurement of poverty in Latin America, describes the difficulties in measuring several components of household income including production for own consumption at the lower end of the distribution, and capital income at the upper end of the distribution. Data collection is improving over time, but he believes that confidence intervals around estimates of the Gini coefficient should be of the order of 4 or 5 percentage points. These limitations in the data, he argues, make it hard to determine the impact of the policy reforms in Latin America on poverty. K. Sundaram and Suresh Tendulkar, Ch. 14 in this volume, discuss further problems in the collection of household survey data in India, including reference periods that have changed, and their implications for estimates of poverty. On the basis of a close analysis of the survey data they conclude that, despite numerous sources of non-comparability between the surveys over time, the finding of declining poverty in India in the 1990s is robust.

Data coverage and comparability

As we have seen, all estimates of global poverty rely on household surveys to provide income or consumption distributions within countries. But while the availability of survey data has improved substantially in recent years, the coverage of countries remains a problem. In particular, the coverage of countries in sub-Saharan Africa is a major concern.

Chen and Ravallion's (2008) estimates of global poverty for the World Bank are based on 675 surveys across 116 countries. Ninety per cent of the population of the developing world is represented by a survey within two years of 2005 (Chen and Ravallion, 2008, pp. 13–16). But this includes only 71 per cent of the population of sub-Saharan Africa. Moreover, coverage for this region is very low in the 1980s, and Chen and Ravallion (2008, p. 16) report that "our estimates for the early 1980s rely heavily on projections based on distributions around 1990." Sahn and Younger, Ch.15 in this volume, argue that the World Bank's reported figures for consumption poverty are somewhat implausible for a number of African countries, and reflect the difficulties of collecting expenditure data. Chen and Ravallion (2008, p. 16) also note that their survey data for China in the early 1980s "are probably less reliable than later years." The lack of good quality data will contribute to uncertainty in estimates of global poverty.²⁸

Introduction

In addition to coverage, there is also the question of data comparability.²⁹ In some country surveys, incomes are gross of tax and in others net of tax; for some they refer to cash incomes and for others certain items of income in kind are included. The rental value of owner-occupied housing is imputed in some surveys but not in others. In studies that use secondary data for within-country distributions, different distributions may have different population units—individuals or households (sometimes families)—and these units may be ranked in a variety of ways, such as individuals ranked by income received, individuals ranked by household income per capita (or per equivalent adult), households ranked by household income per capita (or per equivalent adult), or households ranked by total household income. The population unit and ranking concept used to construct the distribution can make a huge difference to measured inequality and poverty. For example, Anand (1983) found that the income share of the lowest 40 per cent varied from 9.6 per cent to 17.7 per cent for differently defined distributions of income from the *same* Malaysian household survey. Chen and Ravallion, on the other hand, use unit record data from household surveys in each country to construct comparable distributions—of individuals ranked by per capita household income (or consumption).

The data comparability issue has been discussed at length by Atkinson and Brandolini (2001) in their review of secondary datasets used in studies of income inequality. On the basis of a detailed analysis of distribution data for OECD countries, they find that problems of comparability, including those described above, are present even in the “high quality” subset of the Deininger and Squire (1996) compilation, used by both Bhalla (2002, and this volume) and Sala-i-Martin (2006). Atkinson and Brandolini (2001, pp. 777–8) conclude that “users could be seriously misled if they simply download the Deininger and Squire ‘accept’ series [i.e. the ‘high quality’ subset].” Recent World Bank distributional data, described in Chen and Ravallion (2004, 2008), are subject to fewer problems of non-comparability than Deininger and Squire (1996)—but some problems remain, such as the unavoidable mixing of income and consumption distributions.

The coverage of PPP data based on the 1993 ICP has also been a significant concern. In the World Bank’s previous global poverty estimates (Chen and Ravallion, 2007), the PPPs for 69 of the 100 countries were based on data collected in the 1993 ICP, while most of the remainder were based on interpolations from cross-country regressions. China and India were important exceptions: India’s PPP was extrapolated from its 1985 estimate, while China’s was “based on a credible independent (non-ICP) study of price levels in ten cities of China” (Chen and Ravallion, 2004, p. 9). As we saw earlier, the 2005 ICP has much wider coverage than the 1993 ICP and includes data for both China and India. The improved coverage has made a substantial difference to estimates of global poverty, which is due in part to the finding that both China

and India are approximately 40 per cent poorer than indicated in previous estimates.

Conclusion

The international community's commitment to halve global poverty by 2015 has been enshrined in the first Millennium Development Goal. How global poverty is measured is a critical element in assessing progress towards this goal. The chapters in this volume address a range of problems in the estimation of global poverty, from a variety of viewpoints. Given their political salience, it is not surprising that controversy surrounds both "official" and independent estimates, and that a lively debate has ensued.

In this introduction, we have discussed different views concerning the possibility of defining and using a meaningful global poverty line. We have examined different PPP exchange rates that have been used to map a global poverty line across countries, and the complications that arise in the periodic re-estimation of PPPs on the basis of different ICPs. One of the most significant differences between the studies is whether they use survey or national accounts means for within-country distributions in the estimation of global poverty. It will be clear from our discussion of the issues that we have reservations about some of the approaches adopted. We have nonetheless tried to include the full range of viewpoints to allow the reader to form his or her own judgment about their relative merits. The debate on the measurement of global poverty will surely continue, and with this volume we hope to illuminate this important topic.

Notes

1. As we discuss below, the original "PPP\$1-a-day" poverty line has been updated to PPP \$1.25 at 2005 prices.
2. Chen and Ravallion (2008, Table 8).
3. Sala-i-Martin (2006, p. 392).
4. PPP exchange rates convert, say, pounds into dollars, not at the official exchange rate but at a rate reflecting the differences in purchasing power. The official exchange rate may be US\$1.50 to UK£1, even if one can buy about the same market basket of goods with a pound in the UK as with a dollar in the US. In this case, the PPP exchange rate would be 1 to 1, not 1.5 to 1. (Not surprisingly, different patterns of consumption will give rise to different PPP exchange rates.)
5. Bhalla (2002) states that "the errors inherent in [National Accounts] means are corrected by increasing the poverty line by approximately 15 percent, from \$1.30 per capita per day to \$1.50 per capita per day" (p. 121). We discuss the use of National Accounts means below.

Introduction

6. More generally, GDP at PPP\$ calculated in year $t+n$ is not equal to GDP at PPP\$ calculated in year t multiplied by intervening domestic growth and deflated by intervening US inflation.
7. The ten countries are: Bangladesh, China, India, Indonesia, Nepal, Pakistan, Tanzania, Thailand, Tunisia, and Zambia (Chen and Ravallion, 2001, p. 285).
8. While Bhalla (2002) and Sala-i-Martin (2006) both use a higher real poverty line than Chen and Ravallion (2007), Sala-i-Martin's estimates of poverty are lower in all years and Bhalla's (2002) are lower from about 1990. This is due to other methodological differences, such as their use of National Accounts data for scaling household incomes, which we discuss later.
9. One could then use the relative cost of this capability set in different countries to infer the implied "PPP" exchange rates. There is no reason to think that such exchange rates would be similar to extant PPP exchange rates.
10. Although it is possible to measure years of schooling reliably, a year of schooling in one country may differ from that in another in quality or in total hours spent at school.
11. Given a benchmark year t , GDP in PPP\$ in year $t+n$ is calculated by scaling GDP in PPP\$ in year t up or down by the country's real growth rate (nominal growth minus a price deflator). In the case of Penn World Tables there is a further stage of reconciliation after this updating (Aten and Heston, Ch. 6). As mentioned in note 6 above, GDP in PPP\$ in year $t+n$ calculated in this manner can be very different from that obtained by use of a new ICP conducted in the year $t+n$.
12. In the case of Bhalla (2002), PWT is only one of several cited sources—see note 15 below.
13. See Anand and Segal (2008, pp. 70–3) for formulas and discussion of the GK and EKS methods.
14. Of course, if the global poverty line itself were re-estimated as the median of the lowest ten poverty lines in GK PPP\$, then the overestimation of incomes in poor countries would also lead to a higher poverty line in GK PPP\$, and such a bias would not be present.
15. Bhalla (2002) mixes PPP sources, which is problematic for a different reason. On p. 140 he reports using the World Bank's *World Development Indicators* as his source for PPPs, but on p. 207 he reports using PPPs additionally from the Penn World Tables, the IMF's *International Financial Statistics*, and Maddison (2001), benchmarked in various years. The PPPs in these sources are inconsistent because some are based on EKS and others on GK, and they are benchmarked to different years (Anand and Segal, 2008, pp. 81–2).
16. It would also require knowing "the poor" in advance, implying a certain circularity. However, this is not a major practical problem.
17. By itself, of course, this does not tell us whether global poverty is under- or over-estimated.
18. All estimates of global poverty are based on datasets that include both income and consumption surveys. The World Bank, wherever possible, chooses consumption surveys over income surveys. Mixing income and consumption surveys raises questions of comparability, which are discussed in Anand and Segal (2008, pp. 73–4).

19. The United Nations University World Institute for Development Economics Research (UNU-WIDER) World Income Inequality Database (WIID) is available at http://www.wider.unu.edu/research/Database/en_GB/database/ (accessed June 13, 2009).
20. Bhalla (2002) also smoothes within-country distributions by fitting a three-parameter Lorenz curve to the quintile shares through regression. After this regression he performs a “filtering” procedure (pp. 133–4), but it is unclear precisely what this procedure involves (Anand and Segal, 2008, p. 81).
21. Bhalla (2002, p. 128) scales to 0.867 times HFCE.
22. The smoothing technique used by Sala-i-Martin (2006) is problematic (see Anand and Segal, 2008, pp. 78–9). First, the theory of kernel density estimation assumes that the income observations are independent identically distributed draws from the underlying income distribution, which is not the case for quintile means. Second, his use of a constant bandwidth across all countries’ datasets is incorrect because the appropriate bandwidth for a distribution depends on its spread. Minoiu and Reddy (2008) find that the use of alternative bandwidths makes a large difference to poverty estimates. Finally, non-parametric kernel density estimation is intended for large datasets, whereas each of Sala-i-Martin’s country distributions is constructed from just five data points (the quintile means).
23. This is, however, somewhat ambiguous. The “Official List of MDG Indicators” on the Millennium Development Goals Indicators website (<http://mdgs.un.org/unsd/mdg/>), accessed June 13, 2009) states the reduction of *income* poverty as the goal, whereas under “Data Availability” the data used to monitor this goal are described as measuring consumption poverty.
24. The latter includes expenditure by organizations such as political parties and religious associations.
25. Aten and Heston (Ch. 6) state that the latest PWT, version 6.1, includes expenditures on health and education by government and non-profit institutions in “Household Actual Final Consumption” for OECD countries, but not for other countries.
26. Much of this paragraph closely follows Deaton (2003, pp. 367–8).
27. See Kulshreshtha and Kar (2005).
28. The coverage and quality of data in Bhalla (2002) are questionable. According to his Table A.1 (p. 209) there are 317 surveys for the period 1950–80, and 604 for the period 1980–2000, implying a total of 921. But he also reports using “more than 1,000 household surveys” (p. 38). Moreover, Ravallion (2002, p. 8) observes that only “[a]bout half of Bhalla’s 600 distributions over 1980–2000 would pass the quality standards applied to the [World] Bank’s calculations.”
29. This paragraph draws on Anand and Segal (2008, p. 74).

Bibliography

- Ackland, R., Dowrick, S., and Freyens, B. (2004) ‘Measuring Global Poverty: Why PPP Methods Matter’. Paper presented at the 2004 Conference of the International Association for Research in Income and Wealth, Cork, Ireland (August).

Introduction

- Ahmad, S. (2003) 'Purchasing Power Parity for International Comparison of Poverty: Sources and Methods'. Mimeo, Development Data Group, World Bank, Washington, DC. <<http://www.worldbank.org/data/ICP>> accessed on January 15, 2007.
- Anand, S. (1983) *Inequality and Poverty in Malaysia: Measurement and Decomposition*. Oxford and New York: Oxford University Press.
- and Segal, P. (2008) 'What Do We Know about Global Income Inequality?' *Journal of Economic Literature*, 46(1), pp. 57–94.
- Atkinson, A. B. and Brandolini, A. (2001) 'Promise and Pitfalls in the Use of "Secondary" Datasets: Income Inequality in OECD Countries'. *Journal of Economic Literature*, 39(3), pp. 771–99.
- Bhalla, S. S. (2002) *Imagine There's No Country: Poverty, Inequality and Growth in the Era of Globalization*. Washington, DC: Institute for International Economics.
- Boskin Commission (1996) 'Toward a More Accurate Measure of the Cost of Living'. Final report to the Senate Committee from the Advisory Commission to study the consumer price index, December 4. <<http://www.ssa.gov/history/reports/boskinrpt.html>> accessed on March 6, 2009.
- Chen, S. and Ravallion, M. (2001) 'How Did the World's Poorest Fare in the 1990s?' *Review of Income and Wealth*, 47, pp. 283–300.
- — (2004) 'How Have the World's Poorest Fared since the Early 1980s?' *World Bank Research Observer*, 19(2), pp. 141–69.
- — (2007) 'Absolute Poverty Measures for the Developing World, 1981–2004'. World Bank Policy Research Working Paper 4211, Washington, DC.
- — (2008) 'The Developing World is Poorer than We Thought, But No Less Successful in the Fight against Poverty'. World Bank Policy Research Working Paper 4703, Washington, DC.
- Deaton, A. S. (2001) 'Counting the World's Poor: Problems and Possible Solutions'. *World Bank Research Observer*, 16(2), pp. 125–47.
- (2003) 'How to Monitor Poverty for the Millennium Development Goals'. *Journal of Human Development*, 4(3), pp. 353–78.
- and Kozel, V. (eds.) (2005) *The Great Indian Poverty Debate*. New Delhi: Macmillan.
- Deininger, K. and Squire, L. (1996) 'A New Data Set Measuring Income Inequality'. *World Bank Economic Review*, 10(3), pp. 565–91. Updated dataset available at <<http://econ.worldbank.org>>.
- Korinek, A., Mistiaen, J. A., and Ravallion, M. (2006) 'Survey Nonresponse and the Distribution of Income'. *Journal of Economic Inequality*, 4(1), pp. 33–55.
- Kulshreshtha, A. C. and Kar, A. (2005) 'Estimates of Food Consumption Expenditure from Household Surveys and National Accounts'. In A. S. Deaton and V. Kozel (eds.) *The Great Indian Poverty Debate*, New Delhi: Macmillan, pp. 102–18.
- Milanovic, B. (2005) *Worlds Apart: Measuring International and Global Inequality*. Princeton, NJ: Princeton University Press.
- Minoiu, C. and Reddy, S. G. (2008) 'Kernel Density Estimation Based on Grouped Data: The Case of Poverty Assessment'. IMF Working Paper WP/08/183, International Monetary Fund, Washington, DC.
- Ravallion, M. (2000) 'Should Poverty Measures be Anchored to the National Accounts?' *Economic and Political Weekly*, August 26 and September 2.

- (2002) 'Have We Already Met the Millennium Development Goal for Poverty?' *Economic and Political Weekly*, November 16.
- (2004) 'Pessimistic on Poverty?' *The Economist*, April 10.
- and Chen, S. (1997) 'What Can New Survey Data Tell Us About Recent Changes in Distribution and Poverty?' *World Bank Economic Review*, 11(2), pp. 357–82.
- Sala-i-Martin, X. (2006) 'The World Distribution of Income: Falling Poverty and . . . Convergence, Period'. *Quarterly Journal of Economics*, 121(2), pp. 351–97.
- Stiglitz, J. E. (1976) 'The Efficiency Wage Hypothesis, Surplus Labor and the Distribution of Income in L.D.C.s'. *Oxford Economic Papers*, 28(2), pp. 185–207.
- United Nations (2000) *We the Peoples: The Role of the United Nations in the 21st Century*. Secretary-General's MDG Report, New York.
- World Bank (1990) *World Development Report 1990: Poverty*. New York: Oxford University Press.
- (2008a) *Global Purchasing Power Parities and Real Expenditures: 2005 International Comparison Program*. Washington, DC.
- (2008b) 'Comparison of New 2005 PPPs with Previous Estimates'. Revised Appendix G to World Bank (2008a), Washington, DC.