Have We Already Met the Millennium Development Goal for Poverty?

By Martin Ravallion

The author is with the World Bank's Development Research Group. These views should not be attributed to the Bank. Helpful comments were received from Nancy Birdsall, Shaohua Chen, Gaurav Datt, Aart Kraay, Branko Milanovic, Dominique van de Walle, and John Williamson.

In a new book, Imagine There's No Country: Poverty, Inequality, and Growth in the Era of Globalization *(published by the Institute for International Economics), Surjit Bhalla purports to overturn prevailing views on how much progress the developing world has been making against absolute poverty. Indeed, Bhalla claims that by 2000 we had already met the Millennium Development Goal of halving the 1990 incidence of extreme poverty—15 years ahead of time. This would be good news, if it was right.*

This paper offers a critical assessment of Bhalla's arguments and evidence.

In *Imagine There's No Country: Poverty, Inequality, and Growth in the Era of Globalization*, Surjit Bhalla challenges the "conventional wisdom" on how much progress the world has been making against absolute poverty. The World Bank's poverty numbers are his main target. The book claims that the Bank has greatly underestimated the extent of poverty reduction over the last 20 years. And Bhalla sees this as deliberate—so that the Bank could attract more resources for fighting poverty. By Bhalla's assessment, economic growth has been far more pro-poor than the Bank's numbers indicate.[1]

Are these claims justified? The next section offers an overview of *Imagine*'s claims. Then the book's measures of poverty are compared with the Bank's, exploring the reasons for the revealed differences. Finally the discussion turns to the implications for our understanding of the association between economic growth and absolute poverty.

"Conventional wisdom" or a straw man?

Imagine oversells itself from the start. The book claims to be concerned with the impact of "globalization" on poverty and inequality in the world. For example, the book's back cover asks: "Who has gained from globalization?" Readers of this book will not, however, come away any wiser about the answer to that question than when they started. *Imagine* does not contain any analysis that could reasonably allow attribution of the measured changes in poverty and inequality to "globalization." We are not given any evidence that would allow one to identify the

[1] I guess it is reassuring that we don't only get criticisms of our poverty counts from the left, who tend to argue the opposite position—that we underestimate the extent of poverty and overestimate its rate of decline; see for example the comments by Robert Wade (2001).

role played by greater openness to external trade (as one aspect of "globalization") in the distributional changes observed, versus other factors such as rising agricultural productivity, demographic factors, changes in the distribution and returns to education and internal policy reforms.

Putting aside *Imagine's* overstated claims about the impact of globalization, the book does offer a serious challenge to what the author sees as the "conventional wisdom" concerning the empirical facts about what has been happening to poverty and inequality in the world. The book identifies three tenets to this perceived conventional wisdom: that inequality has been rising, that there has been mean income divergence across countries, and that poverty has not been falling. *Imagine* concludes that all three are wrong, with the implication that poverty in the world has fallen dramatically over the last 20 years due in large part to economic growth.

On the first, I do not believe there is any "conventional wisdom" that inequality in the world has risen over 1980–2000. Indeed, there are about as many studies showing falling or virtually stagnant inequality amongst people in the world over this period as there are those showing rising inequality. Bourguigon and Morrison (2002) found a small increase in inequality in the 1980s. Milanovic (2002) found a sharp increase in the late 1980s, but a fall in the early 1990s, though with a net rise overall. The standard errors around the Milanovic estimates are so high, however, that (statistically) one could just as well conclude that there has been no change. Chotikapanich, et al. (1997) found a small decrease in the 1980s. Sala-i-Martin (2002) found that the Gini index had no trend either way in the 1980s but fell in the 1990s.

The third tenet is equally questionable as a characterization of the "conventional wisdom." All of the studies I know of indicate that the incidence of poverty fell over the period 1980–2000. Chen and Ravallion (2001) find that the proportion of the population of the developing world living below either \$1 or \$2 per day (at 1993 purchasing power parity) fell between 1987 and 1998, though not enough to entail a drop in the total number of poor, given population growth. All the estimates made by the Bank have indicated a fall in the aggregate poverty rate since around 1980. The only published Bank estimate spanning 1980–2000 is in World Bank (2002), which shows 200 million fewer people poor in 2000 than 1980.

Of course, there has been more progress in some periods than others. Bhalla focuses on the Bank's estimates for 1987 and 1998. He neglects to tell readers that the same paper he cites as the source of the 1987 and 1998 estimates also gives estimates for 1990, 1993, and 1996.[2] That paper shows that the aggregate poverty rate was virtually flat for 1987–93. In that period the two most populous countries, China and India, had sharp growth contractions due to macroeconomic difficulties. But the estimates suggest that poverty reduction resumed after about 1993, with a falling number of poor as well as the proportion of the population.

The second tenet could reasonably be called the "conventional wisdom" amongst economists. Chapter 2 of *Imagine* claims to debunk the view that there has been unconditional divergence in mean incomes across countries. Divergence has become well accepted, with numerous demonstrations of the fact that there has been a positive correlation between initial income and subsequent growth across countries over the post-war period (see, for example, Pritchett 1997). So Bhalla is at odds with quite a few researchers here.

Imagine does not, however, present the statistical tests that one would need to overturn this conventional wisdom. There are tests for convergence in the literature, but they are not found in this book.[3] *Imagine* does present estimates of how the mean income of people in the richest 20 percent of countries in 1960 has evolved since then, relative to the poorest 20 percent. Yes, this suggests signs of catch up. However, there are two points to note about Bhalla's calculation. Firstly, it should have been pointed out that, in terms of per capita income, China is the 19^{th} poorest country in the Penn World Tables in 1960, and clearly this country must figure prominently in *Imagine's* (population-weighted) calculations. One cannot claim that "divergence is a myth" (p.205) based solely on China! It would be of interest to see if Bhalla's conclusion is robust to using the poorest 18 countries instead.

Secondly, Bhalla confuses the between-country component of inequality with mean-income divergence across countries. The former is naturally population weighted, while the latter is not. It is well known that, while there has been mean-income divergence across countries in the post-war period, the between-country component of overall inequality has been falling, for which economic growth in China has been a major contributing factor (see, for example, Schultz

---

[2] The working paper version of the Chen-Ravallion paper cited by Bhalla gave estimates for these intervening years; this version is available at http://econ.worldbank.org/resource.php?type=5. To save space, the final published version (Chen and Ravallion, 2001) only gave the end points.

[3] A standard test is to regress the change in log mean on its initial value though this is probably biased against finding divergence (given measurement error). This can be dealt with using an instrumental variables estimator. One can also study the evolution of the cross-country density of incomes over time. For further discussion of convergence tests see Durlauf and Quah (1999).

1998; Boltho and Toniolo 1999). Bhalla has not refuted divergence; rather he has rediscovered the falling between-country component of world interpersonal inequality.

If there is a contribution in *Imagine* it is not to be found in what the book says about the distributional impacts of "globalization" or these "conventional wisdoms," but rather in its alternative measures of poverty in the world.

Measuring poverty

Table 1 gives *Imagine's* estimates of the headcount index of poverty (percentage of people living in households with an income or consumption per person below the poverty line) for 1990 and 2000. These are compared to the Bank estimates that *Imagine* quotes often for 1987 and 1998, using the most comparable poverty lines.[4] There are differences in the poverty line used, though in both cases the line is designed to have constant real value over time.

Imagine reports that the poverty rate has virtually halved over the decade, implying that the Millennium Development Goal (MDG) of halving the 1990 "$/day" poverty rate by 2015 has already been reached. By contrast, the Bank's estimates indicate a far more modest drop in poverty incidence; if this rate continued then the MDG will not be reached.

Table 1: Alternative estimates of poverty incidence by region

	Imagine		World Bank	
	1990	2000	1987	1998
East Asia	31.3	6.0	26.6	14.7
South Asia	18.5	7.8	44.9	40.0
Sub-Saharan Africa	55.3	54.8	46.6	48.1
Middle East and North Africa	5.2	7.8	4.3	2.1
Latin America	5.3	5.2	15.3	12.1
Eastern Europe	0.0	0.0	0.2	3.8
Developing world	**25.4**	**13.1**	**28.3**	**23.5**

Note: PPP $1.50 per day for *Imagine*; $1.08 for World Bank.
Sources: Bhalla (2002a), Chen and Ravallion (2001).

The most striking difference in the regional composition is for South Asia (SA), which *Imagine* reports to have a poverty rate well below average, while the Bank finds above average poverty. *Imagine* also has a much higher rate of poverty reduction in SA. Another notable

[4] *Imagine* also gives estimates back to 1950; however, the Bank's estimates have not gone back prior to the 1980s on a comparable basis.

difference is for sub-Saharan Africa (SSA); *Imagine*'*s* poverty count is substantially higher than the Bank's for this region. Indeed, *Imagine's* figures for 2000 suggest that the incidence of poverty is seven times higher than the next poorest region, South Asia; by the Bank's estimates it is only 20 percent higher. *Imagine* reports a small drop in the poverty rate for SSA, while the Bank reports a rise over this period. The estimates agree that East Asia had the highest rate of poverty reduction, though the pace was much higher by *Imagine's* reckoning; the book reports an 80 percent fall in the poverty rate for East Asia, as compared to a (still impressive) drop of almost 50 percent reported by the Bank. There are other notable differences for Latin America, the Middle East–North Africa, and Eastern Europe.

How could two estimates of roughly the same thing be so different? Have we really achieved the MDG for poverty, 15 years ahead of time? To answer these questions we have to probe more deeply into the data and methods that underlie these conflicting numbers.

Setting a poverty line

Different people naturally have different ideas of what "poverty" means. This is true between countries as well as within a given country. The 1990 *World Development Report* (WDR, World Bank 1990) explicitly recognized that richer countries tend to have higher poverty lines when converted to a common currency at exchange rates that attempt to assure purchasing power parity (PPP). Amongst poor countries, there is very little income gradient across countries in their poverty lines—absolute consumption needs dominate in a poor country. But the gradient rises as incomes rise, with poverty lines being roughly proportional to mean consumption for rich countries (Ravallion, et al. 1991).

Recognizing this feature of how poverty lines vary, how should we measure poverty in the world as a whole? One might use the poverty lines found in each country. But then one would not be treating people with the same level of consumption the same way, and so the resulting measures would lose meaning as measures of absolute consumption poverty. Relative poverty lines can still be defended if one believes that relative deprivation matters to a person's welfare. For comparison purposes, Chen and Ravallion (2001) provide such relative poverty measures for the developing world and by region.

However, in its efforts at global poverty monitoring, the Bank has taken the position that to measure absolute consumption poverty on a consistent basis across countries one must use a common poverty line. But if one wants a single number for the poverty count in the world, whose poverty line should it be? Since the WDR 1990, the Bank has chosen to measure global poverty by the standards of what poverty means in poor countries, which gave the “$1/day” line. It is fully acknowledged that this is a conservative definition; one could hardly argue that the people in the world who are poor by the standards typical of the poorest countries are not in fact poor.

This poverty line is then converted to local currency using the latest PPP exchange rates for consumption. The best available consumer price indices are then used to convert the international poverty line in local currency to prices prevailing at about the time of the surveys.

Bhalla claims that in its latest update of these estimates, the Bank has lowered the real value of the poverty line (*Imagine*, Chapter 4). Since elsewhere he accuses us of deliberately overestimating the extent of poverty, it is surely odd that we would have wanted to lower the poverty line. The basis of his claim is that the poverty line we use of $1.08 at international purchasing power parity for 1993 used in World Bank (2000) and documented in Chen and Ravallion (2001) has a lower real value in the US than the line of $31 per month at 1985 PPP used by World Bank (1990).

However, as is well recognized by PPP specialists, there is little comparability between the 1985 PPPs used by the 1990 WDR and those now available. The new PPPs are a great improvement in terms of country coverage, though there are continuing concerns about data quality. There is no straightforward way to convert the old $1/day line at 1985 PPP to a new line with base 1993. Instead, the only defensible approach was to go back to the original poverty lines used by the WDR 1990, and recalculate them with the new set of PPPs, and re-estimate the relationship between national poverty lines and mean consumption which led to the original $0.75 and $1/day lines used in the WDR 1990. The former line was that predicted for the poorest country. The $1/day line had been picked by eye-balling the scatter of points in the relationship between national poverty lines and national mean consumption. For the revision we used instead the median of the lowest 10 poverty lines amongst the 33 countries, which gave the figure of $1.08 at 1993 PPP. This is now virtually identical to the poverty line for the poorest country (given how flat the relationship is at low incomes), so there was no longer any distinction between the $0.75 and $1 lines.

Bhalla's preferred approach of simply adjusting the old line upwards for inflation in the US ignores the fact that there has been (in effect) a PPP devaluation in poor countries relative to the US over the period. For example, China's and Indonesia's poverty lines at 1985 PPP are almost identical to their poverty line at 1993 PPP; India's poverty line at 1993 PPP is only 17 percent higher than its poverty line at 1985 PPP. Yet adjusting the 1985 $1/day line for US inflation would entail an upward increase of roughly 50 percent. In other words, if one simply adjusts the $1/day line for inflation in the United States between 1985 and 1993, then one obtains a poverty line that is well above those found in poorest countries. That would entail a recalibration of the ruler.

Note that, with this change in the PPPs, all past estimates were revised by the Bank's researchers as part of the update—back to those done for the 1990 WDR (see, for example, Chen and Ravallion 2001). This was done to try to assure comparability over time in tracking overall progress against poverty.

Using household surveys to measure poverty

Having set a poverty line, the Bank's researchers then count the number of poor, with appropriate weights, from the raw micro data obtained from the survey or from suitable tabulations based on the micro data. This is essentially the method used by the Bank in both its country-level and global poverty monitoring. The poverty lines are applied to distributions of consumption per person (or income if consumption is not available) constructed from nationally representative household surveys. Adjustments to the data are often required for consistency, such as assuring that population weights are used to obtain an unbiased estimate of the individual distribution of household consumption per person. All calculations are done from the primary data (either micro data or appropriate tabulations).

Bhalla depicts the World Bank's use of surveys to measure poverty as its own idiosyncratic invention around 1990 (see Chapters 7 and 13). However, the Bank's methods are no different in this respect than those used by almost all country statistics officers and independent researchers on poverty and inequality at country level. The Bank did not invent this method, and nor did the Bank only start using it in 1990; the Bank's poverty assessments at the country level routinely use this method, both before and after 1990. The Bank is hardly alone in this respect. For example, estimates of poverty in the United States have been based on surveys for decades, as have those for India (with the exception of a period in which a switch was made to the method Bhalla favors—the government of India was severely criticized within India at the time for cooking the books to show an artificial drop in poverty). And just about every other country in the world measures poverty this way. Indeed, the method has not changed in any essential respect for over 100 years, going back to Rowntree's pioneering study of poverty in York England in the late 19th century. If anything, the use of surveys for this purpose has increased over time, as both the availability of good quality national household surveys and the computing power needed to process them have increased. Bhalla's claim that the Bank came up with this method to deliberately overstate the extent of poverty in the world is as preposterous as it is insulting to the Bank's researchers.

Imagine's method is very different. At one level, the method is brilliantly economical with data, in that it generates poverty and inequality measures with wide coverage over space and time from only minimal published secondary sources. Not every country and every year from 1950 to 2000 has a nationally representative survey. Indeed, there are some countries without any, or just one or two over this whole period. Yet in an appendix, Bhalla gives estimates of quintile shares and Gini indices with virtually complete country coverage. This includes countries for which there were no national household surveys until the 1980s or even 1990s; for example, there were no national surveys for China before the 1980s, or Vietnam before the 1990s. Yet the conclusions in *Imagine* are based on estimates of poverty and inequality by year over 50 years for 149 countries.

To get these 7,450 year-country combinations, Bhalla used 921 published distributions from surveys—317 for 1950–80, 604 for 1980–2000. For the former period, this means an average of one survey per seven years per country; for 1980–2000, it is about one survey per four years.[5] The quality of these distributions is questionable, and particularly so as one goes back in time. There are still no well-defined international standards for how the distribution of income by quintiles (say) should be constructed, and there are endless possibilities. For example, the surveys need not be representative nationally and they need not include important components of income in developing countries such as the imputed value of income from own-production activities. The variable used to rank households to assign them into quintiles varies widely, such

[5] Seemingly impressive figures are quoted for "country coverage," such as 89 percent for Latin America over 1950-80. Readers should be warned however that by Bhalla's reckoning a country is deemed to have 100 percent coverage if it has just one survey over this 30-year period.

as whether it is household or per capita income. Bhalla does not appear to have applied any quality filters; anything that looks like a distribution of income qualifies. This has entailed an expansion in the number of distributions included relative to other studies. Sala-i-Martin (2002) used about 500 distributions, drawn from the Deininger and Squire (1998) database. The latter data compilation (also from secondary sources) applies some quality controls, but is still quite heterogeneous, with numerous anomalies (see, for example, Atkinson and Brandolini 2001). For example, for about 40 percent of the distributions found in the Deininger-Squire database, the ranking variable is household income or consumption, not per capita. It also matters a lot whether one uses fractiles of persons or of households, though one cannot even identify which is which in the Deininger-Squire dataset. And these differences can matter to the measures one obtains of poverty and inequality. The fact that household (rather than per capita) distributions were more common in the 1960s and 1970s than 1980s and 1990s also biases comparisons over time, given that household distributions tend to show higher inequality.

The World Bank's database for poverty monitoring has taken a different approach. The compilation used to measure global poverty applies quality filters to the survey data and applies uniform methods of constructing the distributions from the primary data. The surveys must be nationally representative, include a sufficiently comprehensive consumption or income aggregate (including consumption or income from own production) and it must be possible to construct a correctly weighted distribution of consumption or income per person. Currently only 300 surveys meet even these minimal quality criteria (Chen and Ravallion 2001). About half of Bhalla's 600 distributions over 1980–2000 would pass the quality standards applied to the Bank's calculations.

Another difference is that Bhalla has relied entirely on readily available secondary sources, notably quintile shares of incomes or consumptions and national accounts. By contrast the Bank's global poverty aggregates are built up entirely from the primary data—either the raw micro data or special-purpose tabulations designed to meet uniform quality criteria. This means that past errors in estimating distributional statistics can be dealt with.

National accounts vs. surveys

While the aforementioned differences are likely to matter to the estimates obtained, they may well be less important than another difference, namely that *Imagine's* estimates in table 1 entirely ignore the absolute levels of consumption or income found in all these surveys. Instead, Bhalla prefers to measure the average levels of living from the national accounts.

Given the Lorenz curve and a poverty line, it is well known that almost any poverty measure can be derived as a function of the mean of the distribution of income. One can think of the latter as "inequality" as Bhalla does, though this is not strictly correct; the way the poverty measure varies with the Lorenz curve need not accord with any recognizable inequality measure.

Exploiting the mathematical properties of poverty measures, Bhalla combines his own estimate of the Lorenz curve based on published quintile shares with private consumption per capita from the National Accounts (NAS). The latter is deflated by 15 percent across all countries and dates. *Imagine's* method thus assumes that (i) 85 percent of the national accounts consumption aggregate gives the mean consumption or income of households, which then overrides the survey mean, and (ii) that the surveys get inequality right. If one does not accept these assumptions then one cannot accept the numbers obtained. Let us look at these assumptions in turn.

Why would national accounts give the right mean?

Bhalla admits he is using the NAS for the mean as a matter of convenience ("a marriage of convenience" as he puts it on p.120). But he goes further in arguing that it is actually better to obtain mean household consumption or income from the National Accounts rather than the surveys that were designed for that purpose. However, he provides no convincing evidence that the national accounts give a better estimate of the mean income or consumption of households than do the surveys. This is simply assumed.

There are important conceptual differences in the coverage and definitions of the two data sources, which mean that they would not agree even if there were no errors in either the surveys or the national accounts. These differences are well known. The closest item in the NAS to household consumption is what is usually called "private consumption expenditure" (PCE). This typically includes a lot more than household consumption. For example, it also includes the spending of all nonprofit organizations (NGOs, political parties, and so on). Households are essentially "residual claimants" in the national accounts (Ruggles and Ruggles 1986). One first estimates aggregate output; naturally there are measurement issues in doing so. After adding imports, one tries to account for domestic absorption by firms and governments (the increase in inventories held by firms as well as their purchases and those by the government). What is left over is "private consumption", and it lumps the errors in all other components together, with no reason to think that they cancel out.

By contrast, surveys use the information supplied by households. This brings its own problems. There are compliance issues, whereby some of those households sampled do not want to participate or are not home. And those who do comply do not always give accurate answers, either because they do not know or they are deliberately hiding some of their income.

Bhalla acknowledges these differences in a footnote in Chapter 7, but he says that "these differences are well recognized and can be easily removed to obtain NAS estimates of household consumption" (p.104). How exactly are they removed? While the System of National Accounts since 1993 has recognized the difference between household consumption and total consumption, and recommends separating them, this has so far proved impossible for most developing countries.

However, Bhalla has just taken the private consumption numbers from the published national accounts as his data. He adjusts the NAS aggregates 15 percent downwards, by raising his poverty line from $1.30 to $1.50 per day. Why 15 percent? This appears to be little more than a wild guess, stemming from an undefended assumption that only the richest 2 percent of the population are missing from surveys, but are still captured by national accounts (*Imagine*, p.119). No credible basis is given for this 15 percent adjustment or for assuming that the adjustment should be constant over time and across countries. Bhalla has not eliminated the differences with household consumption as he suggests in the aforementioned footnote, because it is virtually impossible to do that.

Imagine's claim that survey means are wildly off the mark is based on the discrepancies he reports between survey-based aggregates and those from national accounts. There is no dispute that surveys typically return estimates of aggregate household income or consumption that fall short of aggregate GDP or consumption from national accounts. However, *Imagine* appears to overstate these discrepancies, at least as they bear on the Bank's poverty estimates. In particular, the ratios of survey aggregates to NAS aggregates that Bhalla comes up with are appreciably lower than those for the World Bank's data base, as used for the estimates in table 1. Elsewhere I have provided a detailed comparison of the survey means from the latter dataset with the private consumption component of domestic absorption in the NAS (Ravallion 2002). It is true that survey aggregates are typically lower than private consumption from the national accounts. For consumption surveys, it is only about a 5 percent difference taking an unweighted average, though there are countries such as India where there is a worryingly large 40 percent difference (Ravallion 2002). For income surveys it is closer to 25 percent on average. (Consumption surveys are known to be preferable on both conceptual and practical grounds.) By contrast, *Imagine's* estimates indicate a 15 percent difference for consumption in 1987 and 35 percent for income surveys. The fall in this ratio over time reported in the book also appears to be larger than in the Bank's data (based on the regressions reported in Ravallion 2002).

Possibly the fact that *Imagine* has applied no quality filters to the underlying survey data has meant that low-quality surveys have been included that leave out key components of income or consumption (such as from own-farm production). Another possibility is hinted at in a footnote to the relevant table in *Imagine* where it is said that: "The figures represent the ratio of the survey mean income (or consumption) with respect to the national accounts GDP per capita (or private final consumption per capita)" (table 7.1, p.109). I take this to mean that GDP has been used for income surveys and PCE for consumption. However, GDP cannot meaningfully be compared with household income as measured in surveys, since GDP includes much more, and cannot be expected to even approximate household income.

The discrepancies between India's National Sample Surveys (as used for measuring poverty) and the national accounts have attracted considerable attention, and understandably so given India's weight in the poverty counts (at least the Bank's). The appearance of sizable divergence between the National Sample Survey (NSS) and NAS consumption aggregates for India in the 1990s is deceptive unless one takes account of the changes in methods used by the NAS. When the central statistics office changed its methods to accord more closely with new international standards set in 1993 there was a sizable upward revision in the NAS consumption estimates. Bhalla appears to have calculated the ratio of aggregate consumption from India's NSS to that in the NAS ignoring the switch in NAS methods; this gives a deceptively large fall in the ratio in the 1990s. On only comparing estimates using the same methods, Sen (2001) finds little sign of a decrease in the ratio of NSS consumption to NAS consumption during the 1990s, though one still finds signs of a longer-term trend decline in this ratio, going back to the 1970s. By unpacking the published aggregates, one can get some idea of the sources of the discrepancies between national accounts and survey aggregates. For India in 1993–94, Sundaram and Tendulkar (2001) have shown that when one reconstructs consumption from the national accounts in a way that is more consistent with the coverage and definition used by NSS, the two data sources converge considerably, though differences still remain.

Why would national accounts give the right poverty count?

For the main purpose at hand of estimating the amount of poverty it is actually immaterial which gives the better estimate of the mean; the important question is which data source gives the better estimate of the poverty measure—that is after all the object of the exercise. Even if one agreed that the national accounts are right for the mean, *Imagine* provides justification for assuming that the error in the surveys is distribution neutral. Underreporting by the rich is thought to be a much more serious problem (in both absolute level and proportionately) than for the poor.[6] It has not been established, and is quite unlikely from what we know, that the discrepancy between these two data sources is entirely due to underestimation of consumption or income levels in the surveys but that they still get inequality right. More plausibly, underestimation of mean income from a survey tends to comes hand-in-hand with an underestimation of the extent of inequality.

The only justification given in *Imagine* for assuming that the discrepancy is distribution-neutral is provided in Chapter 7, which gives estimates of the adjustment factors by decile of consumption per capita that would be needed to reconcile the survey mean consumption for India in 1993–94 with the national accounts. From the description given, the way these calculations have been done will not be clear to most readers, though it is a method that has often been used before. The method entails calculating an adjustment factor specific to each commodity category, assuming that the national accounts aggregate for that category is correct. One then calculates an adjustment factor for each decile, assuming that the rate of survey underestimation is a constant; this assumption is not made explicit in *Imagine*. The only thing driving the results is then the difference in expenditure patterns.

[6] For example, one study found that the mean income of the 10 highest income households in each of 18 surveys for countries in Latin America was generally no more than the average salary of the manager of a medium to large sized firm in that country (Székely and Hilgert 2000).

As it turns out, Bhalla's results using this method do not support his assumption that the discrepancy is distribution-neutral. Indeed, he finds that larger proportionate adjustments are required at higher consumption levels; by Bhalla's calculations, mean consumption of the richest decile should be expanded up by 53 percent, while the figure is 30 percent for the poorest. To get around this difficulty he asserts that the difference in the decile-specific adjustment factors is not so large as to lead one to question his assumption of a constant (distribution-neutral) adjustment. This is entirely unclear, however, since no calculations are given of the impact on measured poverty. And if anything, his method is biased in favor of finding distribution-neutrality. Assuming that the discrepancy is distribution-neutral at the commodity level is hardly a credible basis of testing for the non-neutrality of survey underestimation in the aggregate.

Fortunately, a careful study by Sundaram and Tendulkar (2001) did similar calculations to Bhalla's on the same dataset and found that for categories of consumption accounting for over 75 percent of the consumption of the poor (bottom 30 percent of the population in terms of consumption per person) the divergence between the NAS and NSS estimates was relatively small, and negative in some cases. Sundaram and Tendulkar conclude that a uniform scalar correction of the mean as used by Bhalla would result in a serious overstatement of the consumption expenditure of the poorest 30 percent of the population, and hence produce a spurious reduction in the headcount index. Bhalla references the Sundaram and Tendulkar study in passing, but does not comment on its implications for his study, or explain the apparent differences (at least in interpretation) with his own results.

After noting the discrepancies between survey means and the NAS aggregates, *Imagine* jumps to the conclusion that the poverty line has been drifting up over time in real value. Here and elsewhere in the book, Bhalla asserts that there is an "equivalence" between proportionate changes in the mean and changes in the poverty line (albeit with opposite sign). It is true that the headcount index of poverty is homogeneous in the mean and the poverty line at any given Lorenz curve. (This follows from the fact that the slope of the Lorenz curve at the headcount index equals the ratio of the poverty line to the mean.) However, there is no reason to assume that the Lorenz curve would be unaffected. Indeed, it is hard to believe that the surveys would be getting the mean wrong but inequality right. Our limited knowledge of the problems of under-reporting and noncompliance in surveys does not suggest that the problem is distribution-neutral.

The key point here is that the relevant bias is <u>not</u> in the survey mean, which is not the object we are trying to estimate, but the poverty headcount. Depending on the structure of the errors in the Lorenz curve, the best mean for measuring poverty may well be the survey mean even if one accepted that it was biased downwards (Ravallion, 2000a). When the statistic we are estimating is a function of many interdependent parameters, only correcting one of them does not mean that you have reduced its overall bias. One can think about it this way: for measuring poverty, there is an error in the mean that potentially offsets the error in estimating the distribution around the mean.

Consider the following example. 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 epsilon 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. This also 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 use the Bhalla method 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.

This is just an example. But it serves to illustrate the general point that simply correcting the survey mean need not get you a better measure of poverty, even if you believe that the national accounts give you the correct mean. *Imagine* leaves the key question begging: <u>If you don't believe the overall survey mean, how can you believe the distribution obtained from the same survey</u>?

Mysterious discrepancies

There is also a large discrepancy between the survey-based estimates reported in *Imagine* and those produced by the World Bank. While the Bank estimates 1200 million people living below $1.08 at 1993 PPP in 1998, Bhalla finds only 743 million. It is difficult to figure out what the problem might be from the documentation provided, but there are some telling signs:

- Different PPP rates are being used. Since it began, the Bank's Global Poverty Monitoring Project has used consumption PPP's, which are clearly more appropriate for measuring poverty than PPP's for GDP. The latest update used the Bank's PPP rates for consumption, which were published on our external web site two years ago and so available to Bhalla at the time he was doing his calculations.[7] There is no reason why he could not have used the same PPP rates in attempting to replicate our numbers. Bhalla has also indiscriminately mixed PPP's from different sources. When the PPP's are revised for the Bank's estimates they are revised consistently across the entire data set.

- Different survey means have been used. Although the survey means corresponding to the Bank's global poverty counts have been available for three years on the same web site, Bhalla has used a different data set; I don't know how well the sources match.
- *Imagine's* adjustments for inflation to get to the reference year appear to be different to the Bank's methods. The main clue I found to this is when Bhalla claims that the Bank has underestimated India's mean consumption for the survey year 1993–94 by 9 percent. However, in his own calculations, he has clearly not corrected for inflation between the mid-point of the reference year, calendar year 1993 (the chosen reference year for the International Comparison Project from which the PPPs are obtained), and the mid point of the survey year, as we have done. And the inflation rate was quite high in India during that year. If one adjusts for the 6 months of inflation as one should, one gets our estimate, not *Imagine's*.
- Bhalla claims his methods are "shockingly accurate" (p.214) though very few tests are reported. Nor are any standard errors given for the estimates (though these would be hard to estimate credibly). The author insists repeatedly that readers should be confident of his results being right, but this appears to reflect little more than the author's personal confidence. For example, it is unclear how well he has validated his Lorenz curve estimates. There are many alternative specifications for a Lorenz curve. They all generally give good overall fits, though that can still mean very different estimates of the headcount index of poverty (which is obtained by finding the point of tangency on the Lorenz curve where the slope equals the ratio of the poverty line to the mean). Bhalla uses only one specification, and reports that it fits well for one country, India, and that it has a low average error in predicting the Gini index. However, a Lorenz curve model might come very close for the Gini index, say, but be way off for the poverty rate. In our

[7] See www.worldbank.org/research/povmonitor/.

work we test various specifications, both for fit and for theoretical consistency. The fact that one specific Lorenz curve model gives a good fit for one country does not mean it will fit well for others. Indeed, we find that very different models of the distribution (either Lorenz curves or density estimation) are needed in different countries and even different dates for the same country.

- An Appendix to *Imagine* gives its estimates of the Lorenz shares and Gini index for 1960, 1980 and 2000. There are clearly numerous inconsistencies with our estimates. For example, Bhalla has underestimated China's Gini index for 2000 by nearly 10 percent (over three percentage points). This could have a large impact on the poverty count, depending on where the difference in Lorenz curves is found.
- *Imagine* has clearly not used internally consistent data. As already noted, by aiming for a broad coverage in his distributional data Bhalla has sacrificed data quality. He has pooled distributions that differ in unknown ways in terms of their ranking variable (household or per capita) and in whether they are fractiles of persons or households.

Implications for the debate on poverty and growth

Amongst those who know the literature, I don't think you will find many people who would disagree with Bhalla's conclusion that economic growth reduces absolute poverty. Indeed, we have been pretty confident of this for at least 10 years, based on the available empirical evidence.[8]

However, should one accept the proposition that economic growth reduces poverty based on the evidence in *Imagine*? An "anti-globalizer" who reads this volume with care might justifiably think that *Imagine's* methods have produced this result by construction—that the calculations are rigged to show that growth reduces poverty. In testing whether growth has reduced poverty it is crucial to recognize that there is a potential bias arising from the negative correlation between the measurement errors in the poverty measures and the measurement error in the mean. In Bhalla's case, the errors in the national accounts will automatically be passed onto the poverty measures. When the NAS overestimate (underestimate) mean income they will underestimate (overestimate) the level of poverty, creating a spurious negative correlation

[8] See World Bank (1990, 2000), Ravallion (1995), Bruno, et al. (1998), and Fields (2000).

between growth and poverty reduction. The same thing happens when using surveys for both the poverty measures and the growth rates.

One can go some way toward fixing this problem. One solution is to use the growth rate from one source as the instrumental variable (IV) for that from the other source (Ravallion 2001); this assumes that these errors are uncorrelated. For almost all regions of the world, the NAS growth rate in consumption is a defensible IV for the growth in the survey mean; the exception is Eastern Europe and Central Asia where there is virtually zero correlation between growth rates obtained from surveys and those from NAS (Ravallion 2002). When one regresses changes in survey-based poverty measures on growth in survey means using NAS growth rates as the IV one also finds that growth tends to reduce absolute poverty, though not as much as the uncorrected estimates suggest.

A method that has been used in the past to assess whether growth is "pro-poor" is to compare the actual rate of poverty reduction to that implied by a distribution-neutral growth process. Datt and Ravallion (1992) showed how one can decompose the changes in measured poverty into three components: a contribution due to growth in the mean relative to the poverty line holding the Lorenz curve fixed (at say the initial value), a component due to changes in the Lorenz curve (holding the mean constant) and a component representing the interaction effect between changes in inequality and changes in the mean (Datt and Ravallion 1992). By comparing the actual change in poverty with the change predicted by distribution-neutral growth (the first term in the decomposition) one can see if there has been any "pro-poor" bias in the growth process. For example, this was done in the 1990 WDR (World Bank 1990).

Imagine asserts that the Datt-Ravallion decomposition is wrong. By differentiating the poverty measure as a function of the mean and "inequality," Bhalla writes down an equation for the change in poverty that does not include the third term. (Also see Kakwani 1993.) However, the inherent nonlinearity in how the poverty measure varies with the mean and inequality means that when the differential calculus is applied to that function, as Bhalla does, the result is only exact for infinitely small changes. (And there is certainly no reason why that function should be homogenous of degree one, which Bhalla claims to be true "by definition.") In reality the changes we see are not of course infinitely small; this is allowed for in the Datt-Ravallion method. That is why the growth-redistribution decomposition has its third term. It is not a mistake, as Bhalla claims, but is inherent to the analytics. It only vanishes in the limit as the changes involved approach zero.[9]

Imagine proposes an apparently new method of testing the effect of growth on poverty. To put this in context, it should be noted that a common test when one has multiple observations of spells of growth and spells of poverty change is to regress the changes in poverty on the growth rate (see, for example, Ravallion 1995). This is not an adequate causal model of poverty by any means; for example, there are good reasons to believe that the elasticity of poverty to growth is not constant but varies across a potentially wide range of country circumstances (Ravallion 2001).

Instead, Bhalla regresses rates of poverty reduction (absolute differences or proportionate differences) on the growth rate times a coefficient "gamma" (which he terms the "shape of the distribution elasticity"). Bhalla's "gamma" is the beginning of period product of the poverty line and the density of the distribution at the poverty line. This picks up differences in how concentrated the population is around the poverty line, which clearly matters to the impact of growth.[10]

Bhalla then compares his regression coefficient to those obtained by regressing on the growth rate alone. His coefficient—from what he modestly calls the "correct model"—is higher (more negative). He claims that his "correct model" gives an elasticity of poverty to growth that is significantly higher than the "incorrect model" of all his predecessors. On this basis he concludes that growth is more pro-poor than the literature has claimed, to quote him: "It is this misspecification that has possibly led to the popular conclusion that growth during the golden age of development (1980–2000) was not pro-poor" (*Imagine*, p.163).

However, Bhalla has done nothing more than rescale the regressor, multiplying the growth rate by a number (specific to each data point) that is typically less than unity. That is why his own regression coefficient is higher. But his coefficient is <u>not</u> then the impact of growth on poverty. To get the elasticity of poverty to growth from *Imagine*'s method one would have to go back to the data and divide by his "gamma" to get the elasticity specific to each data point. While

[9] For discussion in the context of other decompositions in the literature see Ravallion (2000b).
[10] Bhalla appears to think he has discovered the fact that the impact of growth on the incidence of poverty depends on the slope of the cumulative distribution of income in a neighborhood of the poverty line. However, this fact is well-recognized in the literature. Indeed, there is a box in Chapter 3 of the 1990 WDR that is essentially the same as *Imagine*'s Chart 10.1 and the text to the WDR box gives *Imagine*'s key equation (10.4) in words. The analytic derivation was well known at the time and well before; indeed, the WDR box did not bother to cite sources.

there is likely to be heterogeneity in the elasticity depending on initial conditions (which I will return to) one cannot conclude from Bhalla's analysis that poverty is any more responsive to growth on average than we had thought previously. *Imagine* is just going around in circles.

"Growth is sufficient" misses the point

While a sizeable body of existing evidence supports this book's conclusion that lower poverty incidence typically comes hand-in-hand with economic growth, it is also known that there is considerable variance in the impact of a given rate of growth on poverty. With 95 percent confidence, a 1 percent rate of growth in consumption per person can give you anything from a 0.5 percent rate of decline in the "$/day" poverty rate to 3.5 percent (Ravallion 2001). The main proximate cause of these differences in the impact of growth on poverty is differences in the initial level of inequality. Even if inequality does not rise with growth, high inequality countries have a harder time reducing poverty through growth (Ravallion 1997). Higher inequality tends to mean that the poor have a lower share of the gains from growth.

The case of India is instructive about the deeper causes of the heterogeneity in the impacts of growth on poverty across different socio-economic settings. Poverty incidence in India has been falling at a trend rate of about one percentage point per year since about 1970, and the country appears now to have returned to this trend decline since the macroeconomic difficulties of the early 1990s (Datt and Ravallion 2002, review the evidence). However, performance has been uneven between states. Some states have been doing far better than others, both in the longer term, and in the wake of economic reforms over the last 10 years.

But the growth rate needed to achieve this trend decline has been rising over time. The absolute elasticities of national poverty incidence with respect to both nonagricultural output per capita and agricultural yields have been declining over time, especially the nonagricultural growth elasticity. Here the geographic composition of India's growth has played an important role: widening regional disparities and limited growth in lagging areas has made the overall growth process less pro-poor over time. By and large, economic growth in India has not occurred in the states where it would have the most impact on poverty nationally.

There is also evidence for India that these differences in the growth elasticity of poverty relate in turn to differences in access to infrastructure and social services (health care and education) that make it harder for poor people to take up the opportunities afforded by aggregate economic growth (Ravallion and Datt 2002.) The India story also points to the importance of a more disaggregated analysis than found in *Imagine*.

This heterogeneity in the impact of growth on poverty contains important clues on what else needs to be done by governments to promote poverty reduction, on top of promoting economic growth. According to Bhalla, "such actions are not needed…Growth is sufficient. Period." (p.206). The basis of this claim is the evidence he presents that poverty reduction has come with growth. While I do not find his evidence convincing, his conclusion that growth tends to reduce absolute poverty is right, as a generalization.

But that misses the point. Those people Bhalla attacks who are saying that growth is not enough are <u>not</u> saying that growth does not reduce poverty. They are saying that combining growth-promoting economic reforms with the right social-sector programs and policies to help assure that the poor can participate fully in the opportunities unleashed by growth will achieve more rapid poverty reduction than would be possible otherwise.

In a paper on the implications of his book for the globalization debate, Bhalla (2002b, p.2) makes the most remarkable claim I have heard in that debate: "Evidence suggests that no one has lost out to globalization in an absolute sense." Even if one accepted every number in *Imagine* one could not possibly accept this claim. Finding no change in aggregate inequality or poverty from repeated cross-sectional data is perfectly consistent with there being large numbers of losers, and gainers, at every level of living. The critics of globalization justifiably point to the losers. The supporters point to the gainers. Surjit Bhalla seems only to see the aggregate impact, and concludes that everyone is a winner. John Lennon's song appears to resonate for Bhalla: *"You may call me a dreamer."*

References

Atkinson, A.B., and A. Brandolini. 2001. Promise and Pitfalls in the Use of "Secondary" Data-Sets: Income Inequality in OECD Countries As a Case Study. *Journal of Economic Literature* 39 (3): 771–99.

Bhalla, Surjit. 2002a. *Imagine There's No Country: Poverty, Inequality, and Growth in the Era of Globalization.* Washington: Institute for International Economics.

Bhalla, Surjit. 2002b. Who Has Lost Our to Globalization? Oxus Research and Investments, New Delhi. Photocopy.

Boltho, Andrea, and Gianni Toniolo. 1999. The assessment: The twentieth century-achievements, failures, lessons. *Oxford Review of Economic Policy* 15 (4): 1–17.

Bourguignon, Francois, and Christian Morrisson. 2002. The size and distribution of income among world citizens, 1820–1990. *American Economic Review*, forthcoming.

Bruno, Michael, Martin Ravallion, and Lyn Squire. 1998. Equity and Growth in Developing Countries: Old and New Perspectives on the Policy Issues. In *Income Distribution and High-Quality Growth,* ed., VitoTanzi and Ke-young Chu. Cambridge, MA: MIT Press.

Chen, Shaohua, and Martin Ravallion. 2001. How Did the World's Poorest Fare in the 1990s? *Review of Income and Wealth* 47: 283–300.

Chotikapanich, D., R. Valenzuela, and D.S.P. Rao. 1997. Global and regional inequality in the distribution of income: estimation with limited and incomplete data. *Empirical Economics* 22: 533–46.

Datt, Gaurav, and Martin Ravallion. 1992. Growth and Redistribution Components of Changes in Poverty: A Decomposition with Application to Brazil and India. *Journal of Development Economics* 38: 275–95.

Datt, Gaurav, and Martin Ravallion. 2002. Has India's Post-Reform Economic Growth Left the Poor Behind. *Journal of Economic Perspectives* 16 (3): 89–108.

Deininger, Klaus, and Lyn Squire. 1998. New Ways of Looking at Old Issues: Inequality and Growth. *Journal of Development Economics* 57 (2): 259–87.

Durlauf, Steven N., and Danny T. Quah. 1999. The new empirics of economic growth. *Handbook of Macroeconomics.* Amsterdam: North-Holland.

Fields, Gary. 2000. *Distribution and Development: A New Look at the Developing World.* Russell Sage Foundation. Cambridge and London: MIT Press.

Kakwani, Nanak. 1993. Poverty and Economic Growth with Application to Côte d'Ivoire. *Review of Income and Wealth* 39: 121–39.

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

Pritchett, Lant. 1997. Divergence, Big Time. *Journal of Economic Perspectives* 11 (3): 3–17.

Ravallion, Martin. 1995. Growth and Poverty: Evidence for Developing Countries in the 1980s. *Economics Letters* 48: 411–17.

Ravallion, Martin. 1997. Can High Inequality Developing Countries Escape Absolute Poverty? *Economics Letters* 56: 51–57.

Ravallion, Martin. 2000a. Should Poverty Measures be Anchored to the National Accounts? *Economic and Political Weekly* 34 (August 26): 3245–252.

Ravallion, Martin. 2000b. On Decomposing Changes in Poverty into Growth and Redistribution Components. *Journal of Quantitative Economics* 16 (1): 105–18.

Ravallion, Martin. 2001. Growth, Inequality, and Poverty: Looking Beyond Averages. *World Development* 29 (11): 1803–815.

Ravallion, Martin. 2002. Measuring Aggregate Welfare in Developing Countries: How Well do National Accounts and Surveys Agree? *Review of Economics and Statistics*, in press.

Ravallion, Martin, and Gaurav Datt. 2002. Why Has Economic Growth Been More Pro-Poor in Some States of India than Others? *Journal of Development Economics* 68: 381–400.

Ravallion, Martin, Gaurav Datt, and Dominique van de Walle. 1991. Quantifying Absolute Poverty in the Developing World. *Review of Income and Wealth* 37: 345–61.

Ruggles, Richard, and Nancy D. Ruggles. 1986. The Integration of Macro and Micro Data for the Household Sector. *Review of Income and Wealth* 32 (3): 245–76.

Sala-i-Martin, Xavier. 2002. The world distribution of income (estimated from individual country distributions). Columbia University. Photocopy.

Schultz, T. P. 1998. Inequality in the distribution of personal income in the world: how it is changing and why. *Journal of Population Economics* 11 (3): 307–44.

Sen, Abhijit. 2001. Estimates of Consumer Expenditure and its Distribution: Statistical Priorities after the NSS 55th Round. *Economic and Political Weekly* 35 (December 16):

4499–518.

Sundaram, K., and Suresh D. Tendulkar. 2001. NAS-NSS Estimates of Private Consumption for Poverty Estimation: A Disaggregated Comparison for 1993–94. *Economic and Political Weekly* (January 13): 119–29.

Székely, Miguel, and Marianne Hilgert. 2000. What's Behind the Inequality We Measure? An Investigation Using Latin American Data. Office of the Chief Economist, Inter-American Development Bank, Washington. Photocopy.

Wade, Robert. 2002. Are Global Poverty and Inequality Getting Worse? *Prospect Magazine*, March.

World Bank. 1990. *World Development Report: Poverty.* New York: Oxford University Press.

World Bank. 2000. *World Development Report: Attacking Poverty.* New York: Oxford University Press.

World Bank. 2002. *Globalization, Growth, and Poverty.* World Bank, Washington.