

MEASURING POVERTY IN A GROWING WORLD (OR MEASURING GROWTH IN A POOR WORLD)

Angus Deaton*

Abstract—The extent to which growth reduces global poverty has been disputed for 30 years. Although there are better data than ever before, controversies are not resolved. A major problem is that consumption measured from household surveys, which is used to measure poverty, grows less rapidly than consumption measured in national accounts, in the world as a whole and in large countries, particularly India, China, and the United States. In consequence, measured poverty has fallen less rapidly than appears warranted by measured growth in poor countries. One plausible cause is that richer households are less likely to participate in surveys. But growth in the national accounts is also upward biased, and consumption in the national accounts contains large and rapidly growing items that are not consumed by the poor and not included in surveys. So it is possible for consumption of the poor to grow less rapidly than national consumption, without any increase in measured inequality. Current statistical procedures in poor countries understate the rate of global poverty reduction, and overstate growth in the world.

I. Introduction

A central issue in the debate about globalization is the extent to which economic growth reduces poverty. When economic growth benefits everyone in equal proportion, the incomes of the poor grow at the same rate as does mean income. The fraction of the population whose incomes are below a fixed poverty line must then decline with growth, although the rate at which it does so depends on the position of the poverty line in the income distribution, with growth in the mean generating more rapid poverty reduction the greater the fraction of the population who are near the poverty line. If economic growth is unequally distributed, the effects of growth on poverty reduction will be less (or more) depending on whether the incomes of the poor grow by less (more) than average. So much, but perhaps not much more, is common ground.

Early debates on growth and poverty, much influenced by Simon Kuznets's (1955) dictum that inequality would increase in the early stages of development, tended to argue that growth did little to reduce poverty. Writing in the

1970s, Chenery et al. (1974), Adelman and Morris (1973), Fishlow (1972), and Bardhan (1973) all argued that economic development either left the poor behind or actually made them worse off; see Cline (1975) for a contemporary survey. Taylor and Bacha (1976) constructed a growth model of "Belindia," a tiny rich Belgium in a huge poor India, as an example of "the unequalizing spiral" that they saw as fitting the stylized facts of development. Ahluwalia, Carter, and Chenery (1979), who were among the first to measure global poverty using now standard methods, argued that the effect of growth was limited both by the relatively low growth of the poorest countries, and by expanding inequality within them. When Fields (1977) argued that in the Brazilian economic miracle of the 1960s the poor had actually done better than average, he was robustly challenged by Ahluwalia et al. (1980), who showed that Fields's conclusions were not warranted by his data, which were consistent with an uninformatively wide range of differential growth rates of incomes of the poor and nonpoor. This was surely the truth of the matter; in 1980, the data were not available to provide a clear answer to the question whether the poor did better, the same, or worse than average during the unprecedentedly high rates of growth in many poor countries in the immediate postwar period. Researchers were forced to rely on a scattering of published distributional measures, whose provenance and reliability were often unclear; and indeed, Kuznets's famous article used distributional data for only three rich countries, with a smaller amount of information for three poor ones.

The paper by Ahluwalia et al. (1980) was an important impetus to the establishment of the Living Standards Measurement Study (LSMS) at the World Bank. The original purpose of the LSMS was to measure the living standards of the poor in a standardized way, to remedy the paucity of distributional data in the third world, and to set up a system of household surveys that would both support and cross-check the national accounts, as well as replicating for living standards measurement what the UN's System of National Accounts (SNA) had done for national income accounts around the world; see for example Pyatt (2003).

Thirty years later, the data situation has been transformed. There are two key innovations. First, internationally comparable national accounts, based on purchasing power parity

Received for publication June 9, 2003. Revision accepted for publication April 7, 2004.

* Woodrow Wilson School, Princeton University.

The *Review of Economics and Statistics* Lecture, presented at Harvard University, April 15, 2003. I am grateful to Daron Acemoglu, Bettina Aten, Barry Bosworth, François Bourguignon, Shaohua Chen, Russel Freeman, Paul Glewwe, Carol Graham, Tom Griffin, Ivo Hovinga, Alan Heston, Michael Kremer, Martin Ravallion, Dani Rodrik, T. N. Srinivasan, Nick Stern, John Williamson, and Jeronimo Zettelmeyer for help in the preparation of this paper, as well as comments on an earlier draft. Anne Case made many invaluable suggestions and corrected several errors. The views expressed here are those of the author alone.

(PPP) exchange rates, allow comparisons of average living standards across countries in a way that is not vitiated by the gross inadequacies of conversions at market exchange rates. Making comparisons in PPP units corrects, or at least diminishes, the gross understatement of living standards in poor countries relative to rich, and removes the spurious component of growth among poor countries that comes from the elimination of those differences with economic development. PPP exchange rates were first used for global poverty estimates by Ahluwalia et al. (1979), and their use is by now almost universal. Second, there has been an extraordinary growth in the number of household surveys available to the research community, including several dozen LSMS surveys. For example, the World Bank's most recent set of poverty calculations use data from 297 surveys from 88 developing countries (Chen & Ravallion, 2001). Deininger and Squire (DS) (1996) have collected and tabulated data on more than 2,600 Gini coefficients as well as many measures of quintile shares; the WIDER extension includes more than 5,000 Gini coefficients. The unit record data from many household surveys are now routinely available to researchers, including such previously inaccessible troves as nearly 20 years of data from the Indian National Sample Surveys back to the early 1980s. Notable by its exclusion is any similar access to Chinese official surveys.

Yet the controversies are no more settled than they were 30 years ago, although there is certainly more common ground among economists than there is in the world at large. The professional consensus, based on the DS data and on work by them and many others, is that, contrary to Kuznets's hypothesis, and contrary to beliefs in the 1970s, there is no *general* relationship between inequality and growth, and certainly not one in which growth systematically widens inequality, as would be the case of growth left the poor behind. From this, two important propositions follow. First, *at least on average* (and much depends on whether we are averaging over countries or people), growth is good for the poor (Dollar & Kraay, 2002; Ravallion, 2001), as is the growth that is arguably generated by greater openness (Berg & Krueger, 2003). Second, and again *on average*, the fraction of people in poverty should decline as if growth were neutrally distributed. In particular, the relatively rapid growth in the developing world from 1980 to 2000 must have brought about a rapid reduction in the fraction of the world's population that is poor. And indeed, calculations using the Penn World Tables combined with inequality measures—the technique first used by Ahluwalia et al. (1979)—show rapid poverty reduction in the 1980s and 1990s; see Bhalla (2002), Sala-i-Martin (2002), and Bourguignon and Morrisson (2002). According to these calculations, not only has the proportion of poor in the world declined, but the decline has been rapid enough to offset population growth, so that the actual number of poor people in the world has fallen. According to Bhalla, the first of the United Nations Millennium Development Goals,

halving the number of people living on less than \$1 a day between 1990 and 2015, had already been met when the goal was announced.

These optimistic calculations are starkly at odds with the World Bank's numbers on global poverty. The World Bank, which is endorsed by the UN as official scorekeeper for the poverty Millennium Development Goal, uses household survey data to measure the living standards of the poor, ignoring national accounts estimates, and its calculations show relatively little poverty reduction in the 1990s. Chen and Ravallion (2001), which provides the details of the Bank's calculations, shows a reduction in the proportion of the poor living on less than \$1 a day from 1987 to 1998 from 28.3% to 23.5%; they argue that this modest reduction comes, not from any expansion in inequality within countries, but from relatively slow growth in mean consumption. Across their 88 countries, the population-weighted rate of growth in mean consumption was only 0.90% from 1987 to 1998, compared with 3.3% growth in real per capita consumption in the Penn World Tables over the same period. These estimates exclude the latest (1999–2000) Indian data, whose inclusion will increase the growth of the survey means over the 1990s. There remains a large gap between, on the one hand, the direct assessment of the growth of consumption of the poor through surveys, and on the other hand, the growth that is implied by the growth in average accompanied by no general increase in inequality.

The plethora of new data has not resolved the controversy, because the new sources are mutually contradictory. According to direct measurement in household surveys, growth among the poor of the world has been sluggish compared with the average growth rates of the countries in which they live. Yet there is no documented increase in inequality that would resolve the discrepancy. If we are to accept the surveys, growth in the world is a good deal slower than we are used to thinking from the national accounts data, and what growth there has been in the latest two decades has made only a modest dent in the level of world poverty. If we accept the national accounts, and do not challenge the conclusion that there is no general increase in inequality nor any correlation between growth and changes in inequality, then official poverty numbers are overstated, and we have already made rapid progress toward reducing poverty in the world. This paper explores these contradictions empirically with an aim to providing a sharper characterization and to advancing some first hypotheses about causes and possible remedies.

A note of caution at the outset. Because countries have vastly different populations, statements about averages often depend sharply on whether or not they are population-weighted. A third of the world's poor live in two countries, India and China, and the global poverty counts are much affected by what happens there. When we are interested in the well-being of the people of the world, and in the effects of statistical practice and statistical discrepancies on global

poverty measurements, we must weight by population. There is no reason to downweight the well-being of a Chinese peasant relative to a Ghanaian cocoa farmer, nor to believe that the world is a better place when an African moves out of poverty and an Indian moves in. However, many of my concerns are about the relationship between measurement and the level of development, in which case the appropriate procedure is to take each statistical system as the unit, and to ignore population sizes. Beyond that, many of the political negotiations about poverty, and about measurement—for example, those in the councils of the United Nations and the World Bank—are carried on nation by nation. In consequence, I shall typically present both weighted and unweighted results.

II. Surveys versus National Accounts: All Countries

In this section, I consider the cross-country and intertemporal relationships between survey and national accounts estimates of consumption expenditure per capita. Many commentators have noted the (sometimes substantial) discrepancies between survey estimates and their national accounts counterparts. As we shall see below, there are also long-standing literatures in India and the United States, not only on level differences, but also on the fact that survey means grow less rapidly than means in the national accounts. My analysis and data overlap with those of Ravallion (2003), whose main concern is with regional and global analyses of the statistical significance of discrepancies in the levels and growth rates of the ratios of survey to national accounts consumption. For consumption surveys, Ravallion comes to the optimistic conclusion that the significant discrepancies can be traced back to the disarray in the statistical systems of the transition economies. The lack of significant differences elsewhere reflects the large cross-country variation in the ratios, as well as the fact that when surveys are not weighted by population, the low and falling ratio in India, where approximately a third of the world's poor live, is lost in the variation of the ratios elsewhere. In consequence, it is possible for the survey-to-national-accounts ratios to be insignificantly different from one even though the surveys and national accounts data have radically different implications for trends in global poverty.

National accounts estimates of consumption are available for most countries in most recent years, so the countries and dates of the comparison are set by the availability of the surveys. The surveys used here come from a convenience sample assembled from various sources. In most cases, I have survey estimates of mean income or mean consumption from the estimates assembled by Chen and Ravallion (2001), and which appear on the World Bank's poverty monitoring Web site. To these I have added my own estimates for India, most of which appear in Deaton and Drèze (2002), a number of OECD surveys, particularly from the Consumers Expenditure Survey and Current Population

TABLE 1.—DESCRIPTION OF SURVEYS USED IN THE ANALYSIS

Year	No. of Surveys	No. of Countries	Population Covered (bn)	Fraction of World Population (%)
1979	3	3	0.35	9.0
1980	7	6	1.33	34.1
1981	5	5	0.27	6.7
1982	3	3	0.33	8.1
1983	2	2	0.97	23.6
1984	7	5	0.48	11.4
1985	13	11	1.59	37.4
1986	21	18	1.75	40.3
1987	23	21	1.82	41.2
1988	26	24	2.92	64.9
1989	31	28	2.14	45.7
1990	23	20	1.69	35.1
1991	29	26	1.95	38.6
1992	37	34	2.16	42.1
1993	45	41	2.61	49.9
1994	35	30	3.22	60.2
1995	51	45	3.67	67.7
1996	48	44	3.82	68.3
1997	43	38	3.38	61.1
1998	57	53	3.86	70.3
1999	22	19	2.16	39.1
2000	26	23	3.54	63.6
All	557	127	—	—

Notes: Surveys are a convenience sample where survey means were readily available. When the number of surveys exceeds the number of countries, some countries have estimates of both mean income and mean consumption per capita. China is included in 1980, 1985, and every year thereafter; India in 1983, 1988, 1994, 1995, 1996, 1997, 1998, and 2000, but not in 1999. (Indian surveys that run from midyear to midyear have been arbitrarily allocated to the second year.) There are 278 estimates of mean consumption, and 281 estimates of mean income.

Surveys in the United States, and the Family Expenditure Survey (now the Expenditure and Food Survey) in the United Kingdom, as well as a number of additional survey estimates supplied by the Bank, but not used in their poverty counts, for example, estimates of mean consumption per head from the official Chinese surveys. In all, I have 557 survey-based estimates of mean consumption per head or mean income per head (occasionally both). Table 1 shows that these come from 127 countries; that the earliest year is 1979 and the latest 2000. The number of surveys in the data set grows steadily larger over time; I have only 3 in 1979 and 7 in 1980, but 57 in 1998 (the peak year). There are 22 surveys for 1999 and 26 in 2000, but this diminution in numbers after 1998 reflects merely the delay in processing and obtaining survey data, rather than any slackening in the growth of usable surveys around the world. For a single country, consumption and income estimates may come from the same survey (for example, in China) or from different surveys (for example in the United States). The fraction of the world's population covered by the surveys shows a strong upward trend, more than doubling from 1980 to 1998, but there is fluctuation in the fraction from year to year as individual countries move in and out of the counts. Much depends on whether or not there is an Indian survey in a specific year. China is included in 1980, and from 1985 onward.

Table 2 shows information on the ratios of survey estimates of consumption or income per head to consumption or income per head from the national accounts. The ratios

TABLE 2.—RATIOS OF SURVEY MEANS TO MEANS FROM NATIONAL INCOME ACCOUNTS

	No. of Surveys	Unweighted			Population-Weighted		
		Mean Ratio	Standard Error	Standard Deviation	Mean Ratio	Standard Error	Standard Deviation
<i>Consumption to Consumption</i>							
All	277	0.860	(0.029)	0.306	0.779	(0.072)	0.191
EAP	42	0.819	(0.069)	0.224	0.863	(0.031)	0.110
EECA	59	0.847	(0.038)	0.230	0.796	(0.040)	0.184
LAC	26	0.767	(0.094)	0.329	0.585	(0.078)	0.193
MENA	20	0.955	(0.104)	0.300	0.867	(0.111)	0.270
OECD	33	0.781	(0.052)	0.097	0.726	(0.032)	0.076
SA	23	0.649	(0.063)	0.122	0.569	(0.036)	0.103
SSA	74	1.000	(0.061)	0.415	1.089	(0.089)	0.459
<i>Income to Consumption</i>							
All	266	0.904	(0.034)	0.290	1.008	(0.044)	0.174
EAP	32	1.036	(0.065)	0.244	1.057	(0.019)	0.105
EECA	47	0.852	(0.038)	0.231	0.811	(0.030)	0.196
LAC	100	0.893	(0.084)	0.392	1.004	(0.143)	0.416
OECD	75	0.891	(0.020)	0.137	0.910	(0.011)	0.084
SA	8	0.892	(0.028)	0.118	0.874	(0.009)	0.101
SSA	4	1.000	(0.136)	0.420	1.023	(0.204)	0.359
<i>Income to GDP</i>							
All	272	0.569	(0.023)	0.203	0.542	(0.023)	0.113
EAP	32	0.515	(0.031)	0.124	0.512	(0.007)	0.051
EECA	49	0.530	(0.029)	0.157	0.481	(0.016)	0.119
LAC	103	0.616	(0.055)	0.264	0.661	(0.104)	0.288
OECD	76	0.527	(0.027)	0.092	0.586	(0.018)	0.059
SA	8	0.685	(0.009)	0.100	0.659	(0.010)	0.071
SSA	4	0.837	(0.138)	0.512	0.672	(0.098)	0.228

Notes: EAP is East Asia and Pacific, EECA is eastern Europe and central Asia, LAC is Latin America and the Caribbean, OECD comprises the countries of the OECD, SA is south Asia, MENA is Middle East and north Africa, and SSA is sub-Saharan Africa. There are no income surveys for MENA in the sample. Numbers differ slightly from table 1 because the relevant national income magnitudes are not always available. Panel 1 shows the ratio of consumption from the survey to consumption from the national accounts, panel 2 the ratio of income from the surveys to consumption from the national accounts, and panel 3 the ratio of income from the surveys to GDP from the national accounts. Standard errors are calculated so as to allow for correlations within countries.

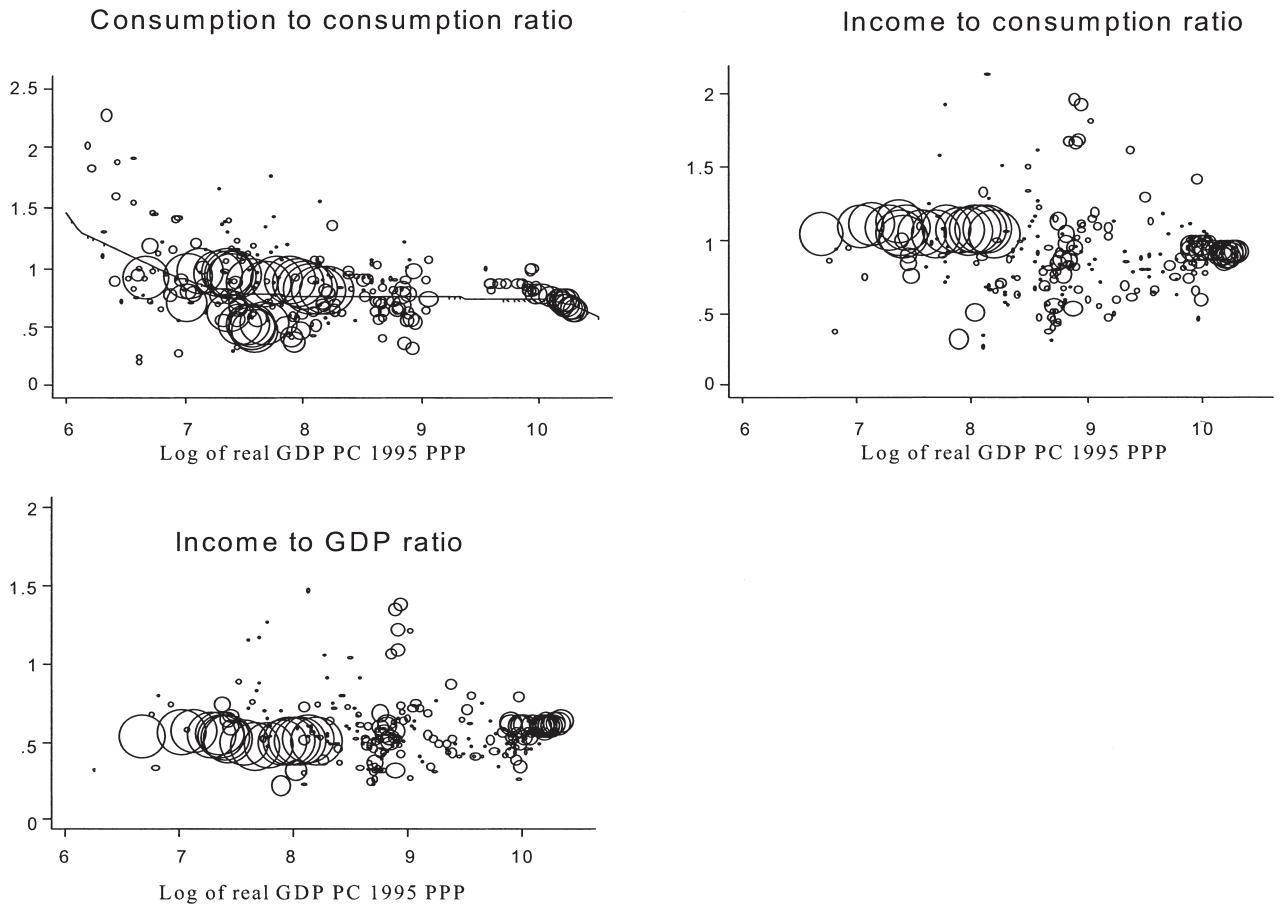
are calculated using nominal values in local currency units (lcu) for both the numerator and denominator. National accounts estimates of household final consumption are the obvious counterparts to survey consumption. For income, most countries do not publish data on disposable household income, so that possible counterparts are GDP or, once again, household consumption. The argument for the latter is that much of saving may not be done by households, but by corporations, government, or foreigners, so that household income may be closer to household consumption than to national income. The top panel shows summary statistics for ratios of survey to national accounts consumption per head, the second panel is for the ratios of survey income to national accounts consumption, and the third panel is for ratios of survey income to GDP.

Consumption estimated from the surveys is typically lower than estimated from the national accounts; the average ratio is 0.860, (with a standard error of 0.029), or 0.779 (0.072) when weighted by population. (India has particularly low ratios.) The exception is sub-Saharan Africa, where the average ratio of survey to national accounts consumption is 1 in the unweighted and greater than 1 in the

weighted calculations. For the OECD, where survey and national accounts quality are presumably the highest, the surveys pick up only a little more than three-quarters of consumption in the national accounts. These differences come in part from differences in definition—for example, national accounts consumption includes such items as the imputed value of owner-occupied housing, which is nearly always excluded from the surveys—but they also reflect errors and omissions in both surveys and national accounts. In consequence, that the ratios for the Middle East and North Africa (MENA) and sub-Saharan Africa (SSA) are close to 1 says nothing about the quality of the surveys in those two regions. Indeed, it is possible that the perfectly measured ratio is less than 1, but is actually measured as greater than 1 because there is understatement in the national accounts. And it is entirely possible that the high ratios for SSA come from large-scale underestimation in the national accounts.

Income measured in the surveys is on average larger than consumption measured in the surveys, but is in most cases less than national accounts consumption, and much less than GDP. Survey income is less than 60% of GDP on average.

FIGURE 1.—RATIO OF SURVEY ESTIMATES OF MEAN INCOME OR CONSUMPTION PER CAPITA TO COMPARABLE NATIONAL ACCOUNTS ESTIMATES



498 Surveys, 124 Countries, Years from 1979 to 2000. The diameter of the circles is proportional to national population in the year of the survey.

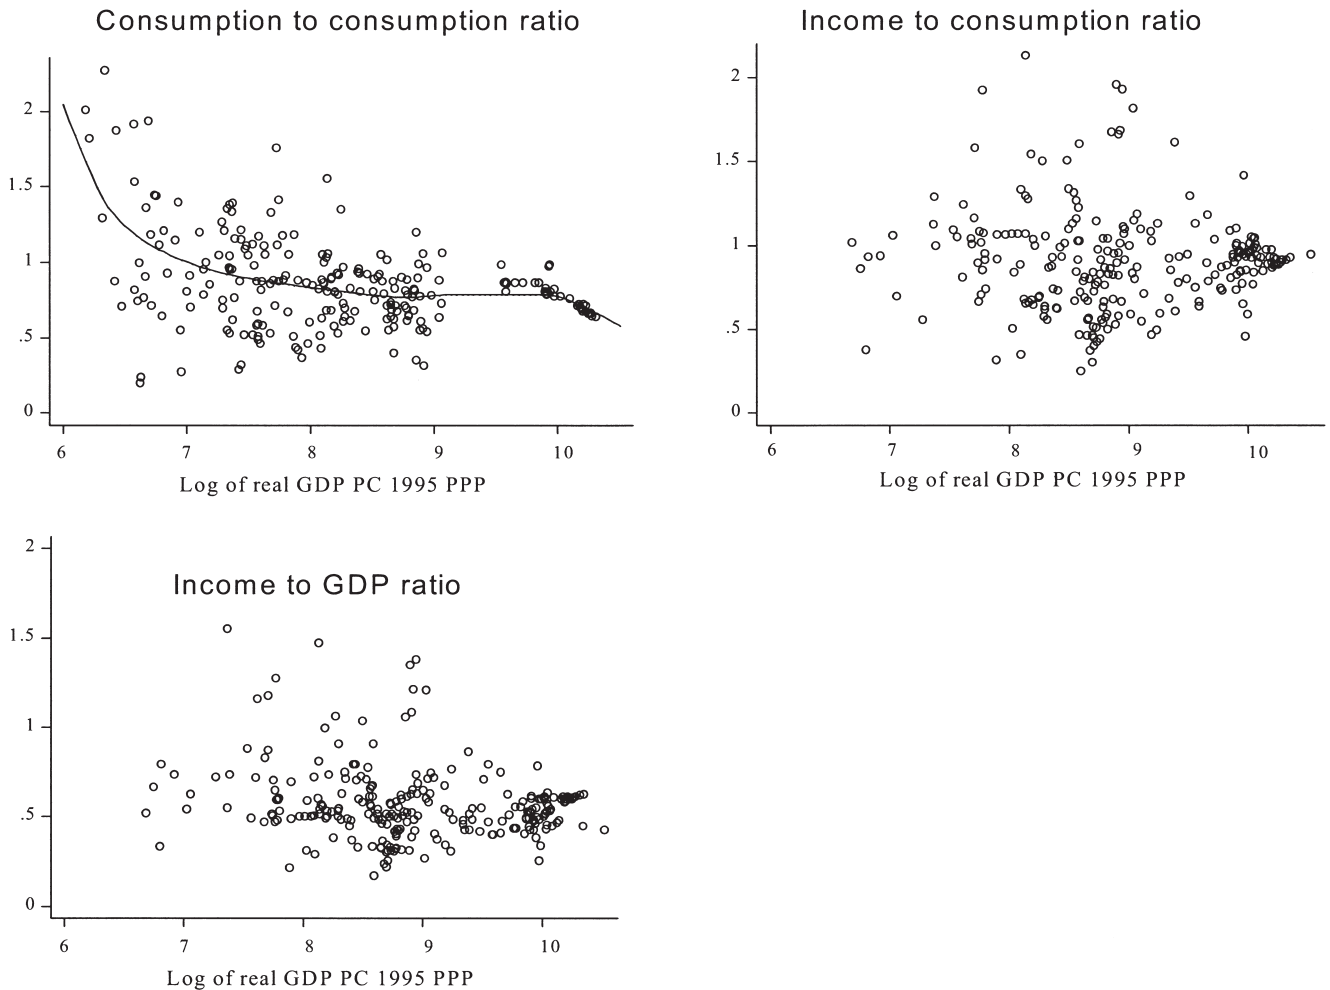
The standard deviations of the ratios provide one crude indicator of combined survey and national accounts accuracy, including both sampling and nonsampling errors. Without prejudging the relative accuracy of national accounts and the surveys, the latter are more likely to vary from year to year, for example because of sampling and changes in survey design, and from country to country, because survey protocols are less standardized internationally than are national accounts. By this measure, the surveys in sub-Saharan Africa are the most problematic, though the surveys in Latin America and the Caribbean also show great variance, particularly the income surveys. OECD surveys have the lowest variance, followed by south Asia, where high-quality household surveys have been in existence for many years. In spite of the difficulties of collecting data in transition economies, the eastern Europe and central Asia region does not show particularly high variance. In several countries in both EECA and LAC, high inflation poses great problems for both survey and national accounts data.

Figures 1 (weighted by population) and 2 (unweighted) show how the same three ratios are related to the level of GDP, here GDP per head in 1995 PPP dollars. (This is the

World Bank's current PPP series, divided by the implicit price deflator of GDP in the United States.) Cross-country and time series data are pooled in these graphs.

There are two points to take away from these figures. First, the top left panels in both figures show a negative relationship between the ratio of survey to national accounts consumption on the one hand, and the GDP per capita on the other. This relationship is steepest among the poorest countries, is flatter in middle income countries, but resumes its downward slope among the rich countries. The continuous lines in the two top left graphs are locally weighted non-parametric regressions of the relationship using a bandwidth of 1.5 (units of real log GDP in PPP). Second, there is no similar relationship among the income surveys, either for the ratio of survey income to national accounts consumption, or for the ratio of survey income to GDP. At least some of the pattern in figure 1 must come from the fact that consumption is typically much easier to measure in surveys than is income in poor countries, where many people are self-employed in agriculture, whereas the opposite is true in rich countries, where most people are wage earners and are more reluctant to cooperate with time-consuming consumption surveys.

FIGURE 2.—RATIO OF SURVEY ESTIMATES OF MEAN INCOME OR CONSUMPTION PER CAPITA TO COMPARABLE NATIONAL ACCOUNTS ESTIMATES




498 surveys, 124 countries, years from 1979 to 2000. Unweighted.

For assessing trends in global poverty and growth, the most important feature of these data is the behavior of the ratios over time. This issue is explored in figure 3 and table 3. Because the subset of countries for which we have survey means differs from year to year, it is not useful to calculate rates of growth of the survey means country by country and then weight by population to obtain estimates of global growth from the surveys. Instead, I have computed population-weighted averages for each year, over whatever subsets of countries have survey data. First, the local-currency consumption and income means are converted to PPP dollars by deflation by the consumption PPP exchange rate from the Penn World Tables, version 6.1, (PWT6.1) and then into real terms by deflation by the U.S. CPI. They are then weighted by population and averaged, excluding the wealthy countries of the OECD. The resulting series are plotted as the bottom two lines in figure 3. They differ in their treatment of country-years where there are both a consumption and an income mean. For the broken lower line I have chosen the consumption survey whenever both are available, and for the solid upper line I have chosen the income survey. (The

results of choosing income means are almost identical if we take income means for China and consumption means elsewhere.)

For comparison with these survey-based estimates, I have used real consumption from PWT6.1 calculated by applying the consumption share to the chain-weighted GDP series. The top solid line in the figure shows the population weighted average of PWT6.1 consumption for all of the countries that ever appear in the survey data set, excluding only the OECD. The broken line is also a population-weighted average of PWT6.1 consumption, but for each year is averaged only over the countries for which there are survey data. This calculation allows a comparison with the survey calculations in which the two series are affected similarly by the variation that comes from the fact that survey countries (and thus the composition of the sum across the world income distribution) changes from year to year. And indeed, the year-to-year variation in the broken-line version of the PWT6.1 consumption series is highly correlated with both survey measures. Of course, the year-to-year (or cyclical) fluctuations in all the series in figure 3

FIGURE 3.—LOGARITHMS OF POPULATION-WEIGHTED AVERAGES OF CONSUMPTION OR INCOME



Household surveys and Penn World Tables, v. 6.1.

TABLE 3.—POPULATION-WEIGHTED GROWTH RATES, 1990–2000: REAL CONSUMPTION OR REAL INCOME, VARIOUS MEASURES, NON-OECD COUNTRIES

	Growth Rate (%/yr)			
	Surveys with Consumption Preference	Surveys with Income Preference	PWT6.1, Matching Surveys by Year and Country	PWT6.1, all Survey Countries
Regression of log on time	1.9	4.0	3.8	2.8
Average rate of growth	2.3	5.0	4.5	2.8

Notes: Columns 1 and 2 show the growth rates of population-weighted survey means. In column 1, whenever there is both an income and a consumption mean for a country year pair, consumption is used. In column 2, whenever there are two surveys, preference is given to income. In both cases, survey means are converted to a constant-price PPP basis by dividing by the product of the U.S. CPI and the consumption PPP exchange rate from the Penn World Tables, Version 6.1 (PWT6.1). For each year from 1990 to 2000, a population-weighted average of the survey means is calculated; note that these averages involve different countries in different years (see table 1). The growth rates are then calculated in two ways, by regression of the logarithm on a time trend (first row) and by calculating the average change in the logarithm over the period. These can be quite different when the series is noisy, as is the case here, because countries come in and out of the average. Columns 3 and 4 show comparable population-weighted growth rates for real PPP (chain-weighted) consumption from PWT6.1. In column 3, consumption from PWT6.1 is used only for country year pairs for which a survey mean exists; this column therefore shares the variability in columns 2 and 3 that comes from the varying selection of countries. Column 4 shows the population-weighted growth rates for consumption from PWT6.1 using all countries for which there ever exists a survey.

(except for the top line) comes as much from the changing selection of countries with different living standards as it does from any genuine fluctuations in the unobservable survey mean over all countries, so that we can use these series only to examine long-run growth, not differences in growth rates over subperiods.

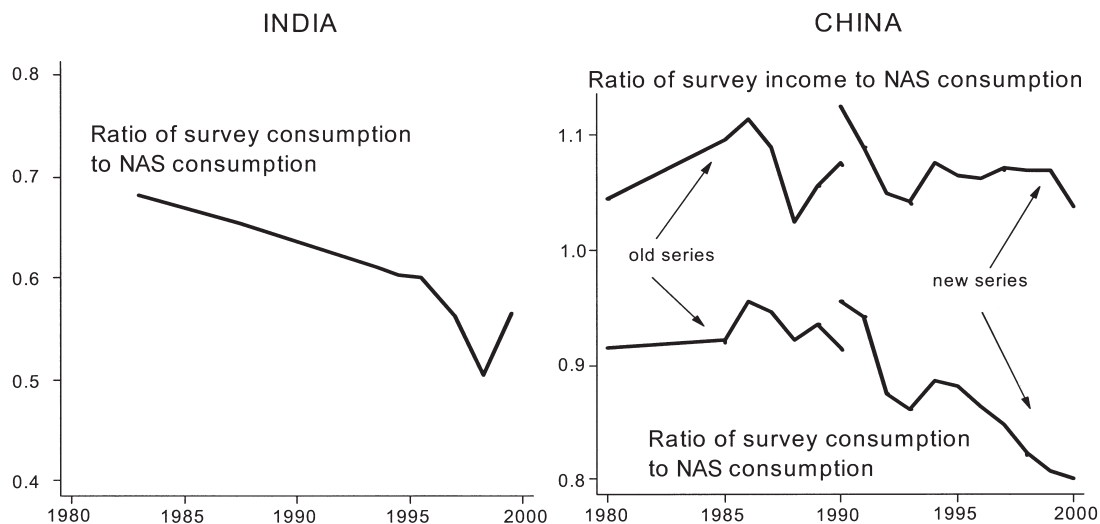
Figure 3 shows that national accounts consumption in non-OECD countries, here taken from the PWT6.1 and shown in the top two lines, grew more rapidly over the 1990s than did consumption from poor countries measured from the surveys, shown in the bottom line. Table 3 shows that growth of survey consumption is 2.3% a year if we

simply take average growth over the decade, or 1.9% a year if we regress its logarithm on a time trend, the difference in the two estimates coming from the variability in the series. This difference is induced by countries with different income levels, particularly India, moving in and out of the survey averages, and is also seen in the comparison growth rates from national accounts consumption, which are 3.8% and 4.5% a year. Whether we take the two low or two high estimates, the growth rate of survey consumption is approximately half of the growth rate of national accounts consumption. If instead of using consumption estimates from the surveys, we take income estimates when they are available, the situation is reversed, and we get a rate of growth from the surveys that is *larger* than the corresponding growth rates in national accounts consumption. The higher growth rate when we give preference to income surveys comes almost entirely from the Chinese data. The World Bank's global poverty estimates use income surveys for China, because there are no distributional data for the Chinese consumption figures. However, in the Bank's calculation the Chinese income distribution is scaled down by the ratio of consumption to income in the Chinese national accounts, a ratio that has been rising over time, so that the first column in table 3 and the bottom graph in figure 3 are the relevant ones for thinking about trends in global poverty as measured by the dollar-a-day counts.

III. Surveys versus National Accounts: India and China

Figure 4 shows the ratios of survey to national accounts estimates for China and for India. The Chinese data, which have a discontinuity in 1990, for which there are two estimates, are from the same survey data base discussed above; the national accounts data are taken from the 2002

FIGURE 4.—RATIOS OF SURVEY MEANS TO NATIONAL ACCOUNTS MEANS OF CONSUMPTION AND/OR INCOME PER HEAD, INDIA AND CHINA




Edition of the World Development Indicators. In China the ratio of survey to national accounts consumption has been declining since around 1990, from a peak of 95% in 1990 to 80% in 2000; the growth rates of the two series thus differ by approximately 1.7% a year in the 1990s. The ratio of survey income (from the same surveys as consumption) to national consumption did not decline over the same period. However, there is a great deal of household saving in China (which shows up in the surveys in that the top line is much higher than the bottom line), so national consumption is not the relevant comparison. Ideally, income should be compared with GDP or, better still, some national accounts estimate of household income. Although I do not have the data to calculate that ratio, there is little doubt that it would also be declining over time. Given the population of China, its increasing discrepancies between survey and national accounts is a major contributor to the global differences. However, it should be noted that many commentators have argued that the growth rates in the Chinese national accounts are too high. The discussions by Maddison (1998), Wu (2000), and Keidel (2001) are all consistent with an overestimation in the rate of GDP growth by between 2 and 4 percentage points a year, and Rawski (2001) argues for much larger overestimation in the last few years. Removing 2 percentage points a year from NAS consumption growth would eliminate the difference in the growth rates between the NAS and the surveys.

Figure 4 also shows the data from India, in this case taken, for national consumption, directly from the latest available edition of the national accounts (Government of India, 2003) and, for the survey estimates, from my own calculations from the unit record data. The Indian National Sample Survey (NSS) conducted its latest full-scale household expenditure survey in 1999–2000, but, because the questionnaire design was changed from earlier similar surveys, there has been controversy about the interpretation of the results. The estimate of average consumption used here

was calculated according to the methods laid out in Deaton (2003) but differs relatively little from the official calculations, much less than is the case for the poverty estimates. In India, survey consumption is much lower relative to national accounts consumption than it was in China. However, as in China, the ratio of the two estimates of consumption has been declining over time. In 1983, the ratio was 0.68, which declined in 1999–2000 to 0.56, so that national accounts consumption has been growing at 1.1%/yr more rapidly than survey consumption. India, like China, accounts for a large share of the world's population, and an even larger share of those who live on less than \$1 a day.

The Indian consumption ratio in figure 4 calls for some additional comment, particularly the erratic behavior from 1995 through 1998. The Indian NSS carries out large household expenditure surveys only once every 6 years or so, the two most recent being in 1993–1994 and 1999–2000. The estimates between those dates come from four smaller NSS surveys that also collect expenditure data. Although the sample sizes of those surveys are sufficient to obtain reliable estimates of the national headcount ratio, there have been questions about their design. The 1998 survey, in particular, lasted only for half a year, and it is arguable that the penultimate observation in the graph should be ignored. Unfortunately, circumstances have conspired to give this and the immediately preceding observation a great deal of weight. Because the 1999–2000 survey was arguably contaminated by changes in the questionnaire, the 1997 and 1998 surveys did not fade into history as quickly as they otherwise would have done. In addition, these were the latest observations for India available to the World Bank for the most recent set of global poverty counts, constructed for the 2000–2001 World Development Report on poverty. The use of the new data in the next round of global poverty counts will give a more optimistic picture of the rate of global poverty decline, though not as optimistic as would be

FIGURE 5.—RATIOS OF SURVEY MEANS TO NATIONAL ACCOUNTS MEANS OF CONSUMPTION AND/OR INCOME PER HEAD: UNITED STATES AND UNITED KINGDOM



the case if survey growth had been as rapid as growth in national accounts.

The internal Indian debate on discrepancies between surveys and national accounts has flared up sporadically for at least thirty years; see in particular the papers in Srinivasan and Bardhan (1974) as well as those of Minhas (1988) and Minhas and Kansal (1989). The recent spate of interest has generated a great deal of important detailed work, including collaborative efforts between the NSS and the National Accounts Division of the Central Statistical Office. Much can be learned from that work, not only for India, but also for other countries, and I postpone discussion until section V.

IV. Surveys versus National Accounts: the United Kingdom and the United States

Although my primary concern is with the measurement of global poverty, and thus with measurement in poor countries, the issue of statistical discrepancies between surveys and national accounts is a general one, and there is a great deal to be learned by looking at the issue at the other end of the global income distribution. Rich countries tend to have fuller data, so that it is sometimes possible to test general hypotheses about surveys that cannot readily be tested in, for example, India or China.

Figure 5 presents the results of survey and national accounts comparisons for the United States and the United Kingdom. The right-hand panel shows results for the Family Expenditure Survey (FES), since 1995 subsumed into the Expenditure and Food Survey (EFS). Data on real consumption per head were taken from the EFS reports, and were scaled up using population and retail price data from the Annual Abstract of Statistics (Office for National Statistics, 2003). The figure shows the ratios of these numbers to final consumers' expenditure from the national accounts. The redesign of the survey and the switch from the FES to the EFS results in a discontinuity at 1995, for which year there are two estimates. The left-hand panel shows corresponding data from the United States using two different surveys, the

Current Population Survey (CPS) and the Consumer Expenditure Survey (CEX). The CEX is the main consumption survey in the United States, although it also collects income data, and the two solid lines in the left-hand panel show (a) the ratio of consumption from the CEX to consumption in the national accounts (the upper solid line) and (b) the ratio of pretax income from the CEX to personal income from the national accounts (the lower solid line). The CEX income and consumption estimates are calculated by the Bureau of Labor Statistics from the CEX, and include estimates from both the diary and interview components of the survey, as well as an estimate of the rental equivalents of owner-occupied homes. The CPS, which is the main income survey in the United States, and which is used by the Bureau of the Census to calculate the official estimates of poverty, does not collect data on consumption. The broken line in the figure is the ratio of income from the CPS to personal income in the national accounts.

Both sets of consumption figures show the now familiar pattern of declining ratios of survey to national accounts consumption. In the United Kingdom, the decline is far from uniform over time, and if the break in the survey in 1995 had been in 1994, it could perhaps have been attributed to the change in design. As it is, the ratio declines by approximately 10 percentage points over the 25 years from 1976 to 2001, so that survey consumption is growing approximately half a percent a year less rapidly than consumption in the national accounts. The decline in the corresponding ratio in the United States is a good deal more dramatic, from 0.80 to 0.64 from 1984 to 2001, so that the difference in the two growth rates is 1.3%/yr, a little more than in India and a little less than in China. Income from the CEX is also a declining fraction of personal income in the national accounts, although the rate of decline is much slower, less than 5 percentage points over 17 years. And income in the CPS shows no trend relative to personal income in the national accounts.

More careful comparisons between the CEX and national accounts consumption data have been made by Triplett

(1997) and by Garner et al. (2003). After making a number of corrections to try to put the two series on a comparable basis, Triplett estimates that from 1984 to 1994, personal consumption expenditures grew at 1.0%/yr more rapidly than consumers expenditure from the CEX. Garner et al. comparing only comparable items, calculate that the ratio of CEX to national accounts values was 89% in 1992. In 1997 and 2000, the comparable ratio was only 80%, so that the differential growth rate was 2.4%/yr until 1997 and 1.5%/yr to 2000.

The differential behavior of income and consumption ratios may have something to do with the fact that in the United States, consumption is much harder to collect than is income. The CEX costs a great deal more per interview than does the CPS, and, whereas for most people (those who are not self-employed) income can be collected with only a few questions, consumption requires a long interview or extensive recordkeeping in diaries. The nonresponse rate in the CEX has been rising over time (Groves & Couper, 1998), while that for the CPS has been constant. It is possible that people are less and less willing to cooperate with the CEX over time, but those who do so are prepared to answer the income questions more fully and more accurately than the consumption questions. The CPS, which does not ask consumption questions, may suffer from fewer problems. Its sample size is also much larger, 60,000 households versus only 5,000 up to 1999, and 7,500 thereafter.

V. Why Do Surveys and National Accounts Diverge?

The previous sections have documented the fact that consumption measured from surveys frequently grows less rapidly than consumption measured from the national accounts. Consistent with this general relationship, the ratio of the two magnitudes is highest in the poorest countries, and lowest in the richest. Within countries as diverse as China, India, the United Kingdom, and the United States, the ratio falls over time as real income increases. Taking non-OECD countries as a whole, population-weighted survey consumption in PPP constant dollars grew at only half the rate of population-weighted consumption in the Penn World Tables. There are conceptual differences between the two concepts of consumption, but these do not account for the differences in growth rates, so that one or both of the growth rates are incorrect. If the surveys are wrong and the national accounts right, either inequality has been widening in ways that our data do not appear to show, or poverty has been falling more rapidly than shown by the \$1-a-day counts. If the surveys are right, there was less growth in the world in the 1990s than usually supposed. Extreme positions apart, we have some combination of underestimation of poverty decline, underestimation of a widening in the distribution of consumption, and overestimation of growth. Quantifying the contribution of each is an urgent task for anyone inter-

ested in growth, poverty, and inequality. In this section, I lay out some of the possible explanations.

It is important to note that there can be no general presumption in favor of one or other of the surveys and the national accounts. In particular, that national accounts are familiar, widely used, and in principle comparable (they typically conform to the UN's System of National Accounts) does not imply that the divergences between them and the surveys must be attributed to the latter. Although it is certainly the case that there exist "failed" surveys, whose execution is known to have been faulty, where fieldwork was disrupted or inadequately supervised, where sampling procedures were flawed, or where changes in survey design made it impossible to compare the results with earlier surveys, national accounts estimates are also subject to many errors, some of which will be discussed below.

A. Unit Nonresponse

Not everyone who is asked to participate in a survey agrees to do so, and failure to respond (unit nonresponse) is known to be different for households with different household characteristics (Groves & Couper, 1998). Of particular interest is the case where better-off households are less likely to respond; Groves and Couper report that, in rich countries, the probability of response is negatively related to almost all measures of socioeconomic status, and though survey organizations in poor countries can usually collect data in very poor areas (albeit under difficult conditions), it is often impossible to penetrate the gated communities in which many rich people live. Suppose then that the probability that consumption y is recorded in the survey is $\pi(y)$, and that $\pi(y)$ is monotone declining in y . This situation has been discussed in a recent paper by Mistiaen and Ravallion (2003), who also show how to use aggregate measures of nonresponse (for example, by region) to correct estimates of poverty and inequality.

If the true (untruncated) density of consumption (or income) is $f(y)$, the density for observed (truncated) consumption is

$$\hat{f}(y) = \frac{\pi(y)f(y)}{\int_{y_0}^{y_1} f(y)\pi(y) dy} = \frac{\pi(y)f(y)}{\bar{\pi}}, \quad (1)$$

where y_0 and y_1 are the bottom and top levels of consumption, and $\bar{\pi}$ is the mean response in the population. From equation (1), the difference in the true and actual densities is

$$\hat{f}(y) - f(y) = \frac{\pi(y) - \bar{\pi}}{\bar{\pi}} f(y), \quad (2)$$

so that the observed density is higher or lower according to whether the household's response rate is below or above the mean. Because $\pi(y)$ is monotone decreasing, the truncated

density is higher at low levels of y , and lower at high values, so that the distribution function is shifted to the left, that is,

$$\hat{F}(y) \geq F(y). \quad (3)$$

This inequality says that the truncated distribution is first-order stochastic dominated by the untruncated distribution, which implies that the estimated poverty rate from the actual data will be no less than the estimated poverty rate in the population, no matter what the poverty line, and that the estimated mean consumption will be no larger than the population mean.

Mistaen and Ravallion also consider the effects of the truncation on the Lorenz curve $L(p)$. In general, the derivative of the Lorenz curve satisfies (see, for example, Kakwani, 1987)

$$L'(p) = \frac{F^{-1}(p)}{\mu} = \frac{y}{\mu}, \quad (4)$$

where y is the p th quantile of consumption and μ is its mean. Provided that $y_0 > 0$ and $y_1 < \infty$, and provided $\pi(y) > 0$, for all y in the support, so that the support of the truncated distribution is identical to that of the original, the reduction in the mean by the nonresponse implies that the truncated Lorenz curve is at least as steep as the true Lorenz curve both at the origin and at $(1,1)$, so that either the Lorenz curves are identical, or they must cross at least once. This result, although obtained under special assumptions (for example, if $y_0 = 0$, it is possible to construct cases where the curves need not cross), tells us that with greater nonresponse by the rich, there can be no general supposition that estimated inequality will be biased either up or down by the selective undersampling of richer households. (The intuition that selective removal of the rich should *reduce* measured inequality, which is sometimes stated as obvious in the literature, is false, perhaps because it takes no account of reduction in the mean from the selection.)

If we are prepared to place restrictions on the compliance function $\pi(y)$, we can analyze the effect of inequality on compliance. In particular, suppose (a) that $\pi(y)$, in addition to being monotone decreasing, is convex, and (b) that $y\pi(y)$ is monotone increasing and concave. Then if F_1 and F_2 are two distributions of income with the same mean, such that F_1 second-order stochastically dominates F_2 , we have

$$\int \pi(y) dF_1(y) \leq \int \pi(y) dF_2(y), \quad (5)$$

so that the average compliance is lower for the more equal distribution. In addition,

$$\int \pi(y)y dF_1(y) \leq \int \pi(y)y dF_2(y), \quad (6)$$

so that, dividing equation (5) by (6), we have

$$\mu_1 \geq \mu_2. \quad (7)$$

Provided the two monotonicity assumptions and convexity and concavity assumptions are satisfied, a mean-preserving increase in spread in the true distribution will decrease the truncated mean. The monotonicity assumption in (b) guarantees that, in spite of the noncompliance, reported income increases with actual income. The concavity and convexity assumptions guarantee the result, but do not appear to be required by the logic of the problem.

To sharpen intuition further, consider the following illustrative but not unrealistic case in which a log normal distribution of income is combined with a probability of compliance that is nonincreasing in income. Suppose that x is the *logarithm* of income or consumption, and that the distribution prior to truncation is log normal with mean (of logs) v and variance (of logs) σ^2 . Suppose too that the probability of responding to the survey is unity up to some income level $\exp(v - \theta\sigma)$, for some number θ , but that above $\exp(v - \theta\sigma)$, the logarithm of the compliance probability declines linearly in the logarithm of income; the kink in the response function is needed to prevent the probability being greater than unity. Hence if $\pi(x)$ is the probability that a household with (log) income x agrees to cooperate, we have

$$\pi(x) = \begin{cases} 1, & x \leq v - \theta\sigma, \\ \exp[-\alpha(x - v + \theta\sigma)], & x \geq v - \theta\sigma, \end{cases} \quad (8)$$

so that the probability of response is unity at the bottom of the distribution. The parameter α is nonnegative, and is (minus) the elasticity of compliance with respect to income.

In the Appendix, I show that, provided θ is large enough, so that noncompliance begins far enough below the mean, the observed (truncated) distribution of incomes is approximately log normal, and that the variance of log income is unchanged, but the mean of logs is shifted downward from v to $v - \alpha\sigma^2$. Although this result is entirely driven by assumption, it illustrates a number of important points. First, we have a case where nonresponse drives the difference between the national accounts and the surveys, and where the mean is biased down, but the Lorenz curve is correct. Second, the ratio of survey consumption to true consumption depends on the variance of the true (and truncated) distribution. In particular, If $\hat{\mu}$ and μ are the truncated and true means of income, the ratio satisfies

$$\ln\left(\frac{\hat{\mu}}{\mu}\right) = -\alpha\sigma^2, \quad (9)$$

so that the understatement of income will be greater in places and at times where inequality is higher. In particular, increasing inequality of incomes will drive down the survey estimates in relation to the truth, even though the ratio of survey to the true mean is independent of the level of mean income. Third, in this case, the ratio of the truncated to the true mean is independent of mean income, so that, although compliance is declining in income, and although average

compliance is declining as the economy expands [at least if the compliance probability in equation (8) is scaled to respond to actual income, rather than the deviation of income from the mean], the fraction of total income captured by the survey does not decrease with growth.

The compliance probability in equation (8) can be generalized, for example by introducing a quadratic term in the second branch of (8), so that

$$\pi(x) = \begin{cases} 1, & x \leq v - \theta\sigma, \\ \exp(-\alpha(x - v + \theta\sigma)) & \\ -\frac{\gamma}{2}(x - v + \theta\sigma)^2, & x \geq v - \theta\sigma. \end{cases} \quad (10)$$

The parameter γ can be positive or negative; in the latter case, equation (10) needs to be modified at high levels of x to stop the probability exceeding 1. Although I do not deal with the complication here, high values of x can be handled in the same way as low values of x in equation (8). Under the same condition as before, that θ is large enough, equation (8) also implies that the truncated distribution will be log normal, but now both mean and variance of logs are changed. Similar algebra to the linear case gives

$$\hat{\sigma}^2 = \frac{\sigma^2}{1 + \gamma\sigma^2} \quad (11)$$

for the variance of logs in the observed distribution, which can be greater than or less than σ^2 , depending on the sign of γ . For the mean of logs, we have

$$\hat{v} = v - \frac{\sigma^2(\alpha + \theta\sigma\gamma)}{1 + \gamma\sigma^2}. \quad (12)$$

Once again, the inequality of income affects the ratio of the observed to the true mean. However, it is no longer appropriate to replace the mean by its true value, leaving the variance unchanged, because if γ is nonzero, the variance is now also affected by the noncompliance, something that we would generally expect to be the case. Note that, as in the original case, the ratio of true to measured income does *not* vary with the true mean, so that noncompliance can increase with income, without the ratio of measured to true income falling with increases in mean income.


There are no ideal aggregate data for testing the extent to which mean income and income inequality affect survey means through noncompliance. Although there is a great deal of distributional information in the DS data set, the information for developing countries is neither reliable in itself, nor well matched to the surveys in the sample discussed above. For the smaller subset of 111 consumption and 77 income surveys for which Gini coefficients are provided on the World Bank's poverty-monitoring web site, there is no significant (unweighted, as is appropriate here) relationship between the log of the ratio of survey to NAS

mean and the Gini coefficient, whether or not real GDP per capita is controlled for. (Region by region, there is a marginally significant effect in South Asia, where the data quality is probably highest; note that the OECD countries are not represented in the poverty monitor countries.) This is also true when the Gini is replaced by the log variance, calculated from the formula for the log standard deviation, $\sigma = \sqrt{2\Phi^{-1}[(g+1)/2]}$, which holds when the distribution is log normal (Aitchison & Brown, 1969).

Another place to look is across the states of India, where there exist state net domestic product data which can be compared with the state means from the household surveys. Again, this comparison is far from ideal: the state domestic product accounts are widely believed to be measured with considerable error; and even without error, the ideal comparison would not be with net domestic product, but with consumption. An offsetting advantage, compared with the international data, is that the state survey means and inequality measures are derived from the same surveys, using identical questionnaires and procedures in each state. It should also be noted that the Indian National Sample Survey Organization consistently maintains that noncompliance is rare, and that numerators make repeated visits until people are available or it is convenient for them. The data from the surveys also carry a notation for whether the household actually surveyed was the one originally intended, or whether it is a substitute for the household targeted for sampling. In the 1999–2000 survey, only 1,200 out of more than 70,000 rural households are listed as substitutes, with 1,900 out of 48,900 urban households. Approximately two-thirds of the substitutions are attributed to the informant being away, and less than a quarter to informants being busy or uncooperative.

The state survey means are well correlated with the state estimates; across the 43rd (1987–1988), 50th (1993–1994), and 55th (1999–2000) rounds of the NSS, and using means for only the 17 largest states, the correlations are 0.88 or higher if Delhi is included, and 0.70 or higher if is excluded. If we use the log of the ratio of survey consumption per head to state net domestic product per head as the left-hand side of equation (6), and the variance of logs from the surveys as the right-hand side, the regression coefficient on the variance of logs is -1.39 ($t = -3.3$) in a pooled regression of the three rounds (54 observations from 18 states in each of the rounds), including round dummies. Taking each round separately, the corresponding coefficients (and t -values) are -0.69 (-0.8), -0.78 (-2.6), and -1.44 (-2.2) for the 43rd, 50th, and 55th rounds, respectively. Figure 6 shows the corresponding plots, with each state identified. Taken literally, these estimates suggest that the elasticity of noncompliance has almost doubled in the 12 years between 1987–1988 and 1999–2000, which is certainly consistent with a fall in the fraction of aggregate consumption captured by the surveys. Inequality, within urban areas and between urban and rural areas, has also

FIGURE 6.—RATIO OF SURVEY MEAN OF CONSUMPTION TO STATE NET DOMESTIC PRODUCT AND VARIANCE OF LOG PER CAPITA TOTAL HOUSEHOLD EXPENDITURE (PCE), INDIA, THREE NSS SURVEY ROUNDS



been rising in India (Deaton & Drèze, 2002), which would again depress the ratio of survey to NAS means. Of course, these results are consistent with a wide range of other possibilities; for example, as suggested by a referee, states with more inequality could have higher savings rates, and in the absence of good capital markets, a higher share of investment and a lower share of consumption in state GDP, or higher inequality, could generate more government expenditure, with similar consequences. So the empirical evidence is weak at best.

These data are not suitable for investigating the important question of whether the ratios are lower when mean consumption is higher. This is because the state net domestic product is used in the calculation of the log ratio, so that to include it in the regression is to guarantee a negative correlation, whether or not one actually exists. Another variable that is plausibly important is the degree of urbanization, if enumerators have greater difficulty contacting or obtaining compliance from urban households. In fact, with the Indian state data, it is difficult to tell the urbanization and inequality explanations apart. Urbanization (the fraction of population in the urban sector) can be used to replace the variance of logs in the regression, with similar *t*-values, and when both urbanization and the variance are entered to-

gether, neither is significantly different from 0. Urbanization and inequality are highly correlated in these data, and we cannot tell whether it is high income that poses the problem for the surveys. With only 18 states, I am almost certainly pushing these data too far. Nevertheless, the question of compliance is central to the analysis of survey versus national accounts, and the Indian experience provides some support for the idea that income-related noncompliance explains some part of the shortfall between the surveys and the national accounts, and perhaps even a part of why the shortfall is increasing.

B. Issues Involving National Accounts

Although noncompliance almost certainly explains at least some of the discrepancies between surveys and national accounts, and although there are other problems with the surveys beyond noncompliance, there are serious quality issues with the national accounts estimates of consumption and GDP. I discuss some of the most important in this subsection.

Discrepancies between survey and national accounts estimates of consumption can arise both through differences in definition, and through differences in the relative success

of meeting those definitions. National accounts statistics are compiled according to protocols laid down in the 1993 version of the Systems of National Accounts, SNA93. The SNA93 establishes what is known as the *production boundary*, which defines what is and is not part of consumption and GDP. The production boundary includes all goods and services that are exchanged, as well as *goods* that are not exchanged, such as food produced for home consumption, but excludes *services* that are not exchanged, such as food preparation, home education of children, or minor home repairs, with the important exception of housing services consumed by owner-occupiers. Because the excluded services tend to be replaced by exchanged services as people become better off and substitute market for home production, the measured growth rate of consumption and GDP will be too high, at least when the true rate of growth is positive. Yet this bias cannot explain any of the growing discrepancy between surveys and national accounts, because the nonexchanged services are not included in the surveys either.

Yet the degree of effective coverage of the nonexchanged items will almost always differ between the surveys and national accounts. Surveys almost never directly collect data on implicit rents for owner-occupiers (other than the maintenance component), though it is sometimes possible to use data from the surveys on housing characteristics to estimate hedonic rental regressions, provided there is a local rental market. But few of the surveys used for poverty analysis contain such estimates, which undoubtedly contributes to the finding that survey-to-NAS consumption ratios are typically less than 1, and if the share of the NAS consumption attributable to rents to owner-occupiers is increasing over time, it will also help explain the increasing divergence. I have not made any general study of the latter, but in India the share appears to be more or less constant over time. In the United States the comparison of the CEX and the NAS in figure 5 includes imputed rents in both numbers, so that this cannot be the source of the increasing discrepancy between them. There are also likely to be differences in coverage of nonexchanged goods. Consumption of own production, gifts, and wages in kind are an important part of the total in poor countries, and many good survey offices devote a great deal of attention to collecting such information. For example, the Indian NSS distinguishes purchases, own production, and gifts for several hundred items. The coverage of nonexchanged goods in the NAS will depend on the methodology employed. Some countries use the survey estimates, but in many and perhaps most cases, consumption is calculated as a residual in a process that begins from production. In principle, this is not a problem, but in many countries it would be extremely optimistic to suppose that the measurement of production accurately captures home production.

Consumption surveys, as opposed to income surveys, are likely to capture a good deal of illegal, or legal but con-

cealed (for example, to avoid taxes or regulation), activities. *Purchasers* of such goods and services, unlike their *producers*, often have no incentives to conceal their transactions, and individuals who have substantial income from sources that they are unlikely to report may nevertheless report the consumption that is financed by that income; see Blades and Roberts (2002), OECD (2002). Because of this, and because many surveys collect comprehensive data on nonexchanged production, it is not surprising that, in some of the poorest countries, consumption measured in the surveys is sometimes much larger than consumption estimated in the national accounts. As nonexchanged production becomes less important with economic development, the effect will wear off, and could thus account for at least some of the decline in the ratio of the two measures. However, as Blades and Roberts emphasize, claims that the existence of various nonobserved activities means that a large share of GDP, as much as a quarter or a third, is missed in the national accounts are almost certainly exaggerated. National income accountants understand the nature of the problems, and although their estimates for the nonobserved economy may not be very accurate, they do not omit it.

In addition to imputed rents of homeowners, there are two other important items of consumption that are included in the NAS, but not in the surveys. One is *financial services indirectly imputed*, or FISIM, which is an estimate of the consumption value of financial intermediation. FISIM is measured as the interest paid to banks and other intermediaries, less interest paid by them. The idea is that interest charged to borrowers contains, in addition to the market rate of interest, a charge for intermediation services to lenders, in addition to the market rate of interest, whereas interest paid to lenders is lower than market, with the difference attributed to financial intermediation services to depositors. The difference between interest paid and interest received is therefore a measure of the value of financial intermediation and, since the 1993 revision of the SNA, has been added to national accounts estimates of household consumption. A similar item is included for risk-bearing services, measured from the profits of insurance companies. In India, the value of FISIM increased from close to 0 in 1983/84 to 2.5% of consumption in 1993–1994 (Kulshreshtha & Kar, 2002), so that this item alone accounts for a quarter of a percentage point per year of the difference in annual growth rates between NAS and survey consumption in India. Note also that, to the extent we are interested in measuring the living standards of the poor, it can reasonably be doubted whether the value of such financial intermediation is relevant. In consequence, even if we accept the argument for the inclusion of FISIM in NAS consumption, neither it nor its rate of growth contributes to the living standards of the poor.

The second potentially important item of consumption included in NAS but not in the surveys is consumption by nonprofit institutions serving households (NPISH), which, in most countries, cannot be separated from household

consumption. It is unclear how large these items are, or whether such expenditures are growing relative to total consumption. In the United Kingdom, NPISH in 2001 was 3.9% of total consumption, almost double the 1970 share of 2.1%. It is possible that NPISH are much more widespread in poorer countries (in India, it is sometimes claimed that there is an NGO in every village), but I know of no data on the subject.

National accounts consumption is typically estimated as a residual using the commodity flow method. Starting from an estimate of domestic production of each commodity, net exports and government consumption are deducted, as are the amounts used in investment and intermediate consumption, with the residual attributed to household (and NPISH) consumption. Many of these calculations are done in physical volumes, so that estimation of consumption in currency units, which is what can be compared with the surveys, requires the use of prices and price indices. There are many opportunities for error along this chain of calculation, and, in general, there is no means (other than surveys) of cross-checking the final answer. The measurement of prices is a survey-based activity with its own sampling and nonsampling errors, and it is sometimes difficult to be sure that prices are those actually paid by consumers. Not surprisingly, the monetary value of NAS estimates of consumption are subject to errors and to occasional large revision. Sundaram and Tendulkar (2003) report that the Indian NAS estimate of consumption of fruits and vegetables in 1993–1994 in nominal rupees more than doubled between the 1998 and 1999 versions of the national accounts. The estimate for clothing fell by approximately a half, and that for rent, fuel, and power rose by more than 40%. Even with some canceling out of pluses and minuses, total consumption was revised upward by 14%, an amount which, if used to calculate poverty rates, would cut the Indian poverty rate by a little less than half.

For food, which is a large share of consumption in poor countries, domestic production is typically estimated by multiplying the acreage of land under cultivation by an estimate of yields per acre. The former comes from a land census or survey, which in many countries is done quite infrequently, whereas the latter comes from crop-cutting surveys, themselves of mixed quality. Data on government consumption are usually relatively accurate, as are imports and exports, which typically are subject to direct monitoring by the government. The same cannot be said for intermediate (business) consumption, which is often assessed by applying various ratios to measured production. These ratios come from enterprise surveys or from input-output tables. Once again, these measures are often outdated. For India, Kulshreshtha and Kar (2002) write that their NAS consumption estimated “depend on an assortment of direct and indirect estimates along with various rates and ratios, some of which are based on the results of studies carried out in the distant past.”

The use of outdated ratios and correction factors is particularly problematic when the economy is growing and its structure changing. Kulshreshtha and Kar, in their detailed commodity-by-commodity comparison of food consumption in the NAS and the NSS in India, find that one of the largest discrepancies is for *vanaspati*, a vegetable cooking oil that is widely used in restaurants. This intermediate use of cooking oil should be deducted in the commodity flow calculations, but in fact there is no such correction in the Indian national accounts. In consequence, and because consumers’ expenditure in restaurants is already included in NAS consumption, restaurant use of *vanaspati* is double-counted in the national accounts. And because consumers switch from domestic cooking to purchased meals as they get better off, the more rapid is the growth of the economy, the larger will the overstatement of consumption become. The *vanaspati* example is an extreme case, in that there is no correction for intermediate business consumption, but the same exaggeration of growth will be generated by the use of outdated rates and ratios to assess intermediate consumption in an economy where growth is reallocating economic activity from own production to the market.

Overstatement of consumption and consumption growth through a failure to capture intermediate consumption will also lead to an overstatement of the level and growth rate of expenditure-based measures of GDP. This exaggeration is in addition to the exaggeration associated with the general movement of activity, such as services, from a nonexchanged to an exchanged basis—for example as a greater share of food preparation is done by food vendors, which is counted in GDP, rather than by family members, which is not. Both come from the same fundamental trend, which is the increasing marketization, complexity, and roundaboutness of production with economic development. Note that not all of these errors in constructing consumption necessarily find their way into GDP. For example, how a commodity flow is allocated between consumption and capital formation will affect the estimation of both, but not of their sum.

C. Other Survey Issues

The two previous sections have documented what are perhaps the most likely candidates for explaining the divergence between national accounts and survey-based estimates of consumption. However, it should also be emphasized that there are many other problems, some of which are on the survey side. It is clear that details of survey design matter for the results, and that protocols are not the same across countries, or sometimes within countries over time. Many of these are discussed in more detail in Deaton and Grosh (2000).

Surveys often have less than complete *coverage*, excluding for example students, the military, and institutionalized persons, expenditures by whom are included in NAS estimates of consumption. In some cases, survey coverage

excluded rural households, or parts of the country that are expensive or dangerous to visit.

Survey questionnaires differ in the length of the *recall period* over which respondents are asked to report their consumption. The choice of recall period is often thought to involve a tradeoff between accuracy of memory, which calls for a short period, and the match between consumption and purchases, which is more accurate when averaged over a long period. But there is little understanding of the effects of different recall periods, particularly in poor, agricultural societies. In India between 1989 and 1998, the NSS experimented with different recall periods, replacing the traditional 30 day recall period for all goods with a 7 day recall period for food and tobacco, and with a 365 day period for durable goods and some other infrequently purchased items. The sample was randomly divided, and half were given the old questionnaire and half the new, so that it is possible to make a clean evaluation of the effects of the change. The shorter reporting period increased reported expenditures on food by around 30%, and total consumption by about 17%, very much in the right direction to help resolve the discrepancy with the NAS. Because there are many Indians close to the poverty line, the 17% increase was enough to reduce the measured headcount ratio by a half, removing almost 200 million people from poverty. What might seem to be an obscure technical issue of survey design can have a major effect on the measurement of poverty, not only in India, but in the world. It should be noted, however, that the higher consumption totals associated with the shorter recall period, although closer to the NAS estimates, are not necessarily more accurate. Indeed, the NSS has carried out a series of controlled experiments in which, for many foods, the 30 day reference period appears to be more accurate than the 7 day period; see NSSO Expert Group on Sampling Errors (2003).

Survey questionnaires also vary in the number of items that are separately distinguished, and there is some evidence that the greater the degree of disaggregation, the greater is the measured consumption in total. There is also no consistency in the treatment of seasonality: some surveys visit each household on several occasions throughout the year, but most simply rely on spreading data collection throughout a calendar year, a procedure that should not bias the mean, though there will be biases in higher-order statistics. In some surveys, respondents keep diaries of their purchases over a period, in others, they make oral responses to interviewers based on recall. Surveys vary on who is chosen as respondent, and whether one or more household members are interviewed. However well informed is the household member who reports purchases, *proxy* reports on the purchases of other family members are likely to be less accurate than reports about the respondent's own behavior. Indeed, proxy reporting can plausibly contribute to a progressively large share of consumption being missed over time. In a poor, rural community, where everyone eats from the same pot and food is nearly all of the budget, the

housewife's report will be quite accurate. This is much less so in more diverse and better-off households, with some family members working outside of the home and maintaining partial budgetary independence.

In addition to the *unit* nonresponse discussed in section V A there is *item* nonresponse, where household members fail to report at least some expenditures, or provide deliberately misleading reports, for example on alcohol consumption or on various illegal items. Finally, and in parallel with the national accounts, there are difficulties in finding adequate prices for consumption items that are not purchased in the market; some surveys use market prices to impute home production, some use farm-gate prices, and some use valuation techniques that are not clearly documented.

It would be desirable if the international statistical community could agree on a common set of best-practice protocols for household income expenditure surveys, as a parallel to the SNA for the national accounts. Unfortunately, most of the problems discussed here are not sufficiently well researched or understood to admit of uncontroversial solutions, and many statistical offices are stout defenders of their own particular practices. Yet, as I shall argue in the next section, only household surveys allow us to measure poverty, so that the task of harmonization must be undertaken if we are to put global poverty measurement on a sound basis.

VI. Conclusions and Implications for the Measurement of Poverty

The standard measures of poverty are based on counting the number of people who live in households whose measured per capita consumption is less than a poverty line. When rich households are less likely to cooperate with the survey than poor people, survey-based estimates of consumption will understate mean consumption and overstate the fraction of people in poverty. Under some conditions, the amount by which average consumption is understated will be larger the greater is the inequality of the true distribution of consumption. Unless consumption inequality is increasing over time, or the fraction of noncooperating households is increasing, income-based noncooperation does not, in and of itself, imply that ratio of measured to true consumption is increasing over time.

National accounts estimates of consumption are typically, although not always, larger than survey-based estimates, and there is a tendency, both across countries and over time within important countries, for the NAS estimate of consumption to grow more rapidly than does the survey-based estimate. Survey-based estimates are subject to numerous errors and inaccuracies, but there are also problems with national accounts estimates. These are likely to understate consumption in the poorest countries, and to overstate the rate of growth of average consumption, both over time in poor countries, and in comparisons between poor and rich countries at a moment in time. In part, these systematic problems in measuring the rate of growth of consumption

carry through to GDP, whose growth rate is also systematically biased upward. I know of no plausible estimates of the size of the bias.

Given the conflict between survey and NAS estimates of consumption, it is tempting to allow the NAS estimates to play at least some role in poverty measurement, instead of using only the survey data. Indeed, the combination of means from the national accounts and Lorenz curves from the surveys has a long history, including Ahluwalia et al. (1979), the Indian government prior to 1993, and most of Latin America until today. In some cases, this procedure was adopted because the survey means were unavailable, and in others, such as the Indian case, the practice was abandoned after searching criticism of the quality of the national accounts; see in particular Minhas (1988), who memorably describes the earlier practice as “mindless tinkering.” In general, there is argument on averaging of multiple estimates (although only estimates of the same thing), so that extensive prior adjustment of any NAS mean would be required before using it to scale up survey estimates. However, there is need for a good deal of caution, and mechanical use of *unadjusted* NAS means, combined with survey-based estimates of the distribution around the mean, will certainly give the poor measures of poverty. There are at least three reasons why.

First, and most generally, the national accounts are designed to generate estimates of macroeconomic aggregates, not estimates of poverty, and the SNA rules are designed with that in mind. National accounts track money, not people. To take an example, the SNA recognizes that production for own consumption is difficult to measure, and recommends that the effort be made only “when the amount produced is likely to be quantitatively important in relation to the total supply of the good in the country” (OECD, 2002, p. 179). Such a rule makes little sense when our prime objective is to measure poverty. At the other end of the spectrum, items like FISIM and the rental value of owner-occupier homes are (properly) included, although in most cases they are either not consumed by the poor, or make up less of their budgets. In general, the NAS is more likely to capture larger transactions than smaller ones, which is close to the opposite of what happens in the surveys, where large *transactors* are the least likely to be included.

Second, the differences in coverage and definition between NAS and surveys mean that, even if everything were perfectly measured, it would be incorrect to apply inequality or distributional measures, which are derived from surveys that measure one thing, to means that are derived from the national accounts, which measure another. When national accounts and surveys are measuring different things, it is perfectly possible for the poor to do less well than the average, without any increase in measured inequality.

To illustrate, suppose that we are interested in measuring consumption growth among the bottom 100p% of the population. We have data on mean consumption μ from the

national accounts, and data on the share s_p of the bottom 100p% of households from surveys. Mean consumption of the bottom 100p% is then estimated to be

$$\mu_p = s_p \mu / p, \quad (13)$$

so that the growth rate of consumption for the bottom group is

$$\frac{\dot{\mu}_p}{\mu_p} = \frac{\dot{s}_p}{s_p} + \frac{\dot{\mu}}{\mu}. \quad (14)$$

In equation (14), the first term on the right-hand side comes entirely from the surveys, and the second entirely from the national accounts. The survey mean $\hat{\mu}$ and the direct survey measure of $\hat{\mu}_p = s_p \hat{\mu}$, the average consumption of the bottom 100p%, are discarded, even though the poor rarely refuse to respond, and provide accurate estimates of their consumption. Moreover, the validity of equations (13) and (14) depends on being able to apply the survey shares to the NAS means, which assumes not only that the NAS means are perfectly measured, but that both are measuring the same thing. So even if we were to accept that NAS consumption is the concept that we want, and even if we were to believe that it is accurately measured, the shares from the survey would be shares of consumption excluding consumption on rents of owner-occupiers, excluding FISIM and the profits of insurance companies, and excluding the expenditures of NPISH. Using the survey shares to allocate NAS consumption to the poor and nonpoor assumes that these items are distributed between poor and nonpoor in the same way as are the goods measured in the survey, an assumption that is not true.

Third, we must recognize that neither the mean consumption nor its distribution are accurately measured, either in the surveys or in the NAS. A particular difficulty comes from the mechanical use of the distributional shares and Gini coefficients that come from the DS and WIDER compilations. (Shares can be calculated from Gini coefficients if a particular distribution is assumed, for example, the log normal.) For most poor countries, these measures are of dubious quality, as indeed is recognized by DS. And neither DS nor the WIDER compilation provides the information that would be required to make an informed judgment on the way their numbers were calculated. So if equation (14) is used to construct $\dot{\mu}_p/\mu_p$, and the measures of s_p are noisy, a regression of $\dot{\mu}_p/\mu_p$ on $\dot{\mu}/\mu$ will have a coefficient that is close to 1, essentially by construction, and the worse is the measurement error, the closer the estimate will be to 1. So there is no credibility to the claim that globalization has been good for the poor based on a calculation that applies badly measured distributional shares to (upward-biased) measures of growth from the national accounts. The globalization debate is serious enough that we must genuinely measure the living standards of the poor, not simply assume them. We cannot prove that growth trickles down by

assuming that growth trickles down, nor argue that globalization has reduced poverty without measuring the living standards of the poor.

If the task were the purely statistical one of estimating mean consumption, there would be much to be said for using the average of the mean consumption from the surveys and that from the adjusted national accounts (Deaton, 2001). But if we need to measure poverty in a way that will convince those who are skeptical of the idea that average growth reaches the poor, there is little choice but to use the surveys. This argument is reinforced by the fact that for many countries it is impossible to adjust NAS consumption estimates to make them comparable with the survey totals. None of this says that the surveys are correct, nor that current measures of global poverty are doing a good job of measuring the trends. And because not every country has a survey in every year, they are clearly unsuitable for measuring year-to-year variations (see figure 3). There is too much incompatibility in survey design across countries. The downward bias in survey measures of consumption almost certainly biases upward the World Bank's global poverty estimates, and in as much as it is unlikely that all of the growth discrepancy between the surveys and the NAS is due to faults in the latter, the rate of poverty decline is likely downward biased. We need an international initiative to provide a set of consistent international protocols for survey design, as well as deeper study of the effects of nonsampling errors, particularly noncompliance.

REFERENCES

- Adelman, Irma, and Cynthia Taft Morris, *Economic Growth and Social Equity in Developing Countries* (Stanford, CA: Stanford University Press, 1973).
- Ahluwalia, Montek S., Nicholas G. Carter, and Hollis Chenery, "Growth and Poverty in Developing Countries," *Journal of Development Economics*, 6 (1979), 299–341.
- Ahluwalia, Montek S., John H. Duloy, Graham Pyatt, and T. N. Srinivasan, "Who Benefits from Economic Development? Comment," *American Economic Review*, 70: 1 (1980), 242–245.
- Aitchison, J., and Alan Brown, *The Lognormal Distribution* (Cambridge, U.K.: Cambridge University Press, 1969).
- Bardhan, Pranab K., "On the Incidence of Poverty in Rural India in the Sixties," *Economic and Political Weekly*, 8: February special issue (1973), 245–254.
- Berg, Andrew, and Anne Krueger, "Trade, Growth, and Poverty: A Selective Survey," IMF working paper no. 03/30 (2003).
- Bhalla, Surjit S., *Imagine There Is No Country: Poverty, Inequality, and Growth in the Era of Globalization* (Washington, DC: Institute for International Economics, 2002).
- Blades, Derek, and David Roberts, "Measuring the Non-observed Economy," *OECD Statistics Brief*, no. 5 (November 2002), 1–8.
- Bourguignon, François, and Christian Morrisson, "Inequality among World Citizens: 1820–1992," *American Economic Review*, 92: 4 (2002) 727–744.
- Chen, Shaohua, and Martin Ravallion, "How Well Did the World's Poorest Fare in the 1990s?" *Review of Income and Wealth*, 47: 3 (2001), 283–300.
- Chenery, Hollis, Montek S. Ahluwalia, Clive L. G. Bell, John H. Duloy, and Richard Jolly, *Redistribution with Growth* (Oxford: Oxford University Press for the World Bank, 1974).
- Cline, William R., "Distribution and Development: A Survey of the Literature," *Journal of Development Economics*, 1: 4 (1975), 359–400.
- Deaton, Angus, 2001, "Counting the World's Poor: Problems and Possible Solutions," *World Bank Research Observer*, 16: 2 (Fall 2001) 125–147.
- "Adjusted Indian Poverty Estimates for 1999–2000," *Economic and Political Weekly*, (25 January 2003), 322–326.
- Deaton, Angus, and Margaret Grosh, "Consumption" (vol. 1, pp. 91–133), in M. Grosh and P. Glewwe (Eds.), *Designing Household Questionnaires for Developing Countries: Lessons from Fifteen Years of the Living Standard Measurement Study* (Washington, DC: World Bank, 2000).
- Deaton, Angus, and Jean Drèze, "Poverty and Inequality in India, a Reexamination," *Economic and Political Weekly* (September 7, 2002), 3729–3748.
- Deininger, Klaus, and Lyn Squire, "A New Data Set Measuring Income Inequality," *World Bank Economic Review*, 10 (1996), 565–591.
- Dollar, David, and Aart Kraay, "Growth is Good for the Poor," *Journal of Economic Growth*, 7 (2002), 195–225.
- Fields, Gary S., "Who Benefits from Economic Development? A Reexamination of Brazilian Growth in the 1960s," *American Economic Review*, 67: 4 (1977), 570–582.
- Fishlow, Albert, "Brazilian Size Distribution of Income," *American Economic Review*, 62: 1/2 (1972), 391–402.
- Garner, Thesia I., George Janini, William Passero, Laura Paszkiewicz, and Mork Vandemia, "The Consumer Expenditure Survey in Comparison: Focus on Personal Consumption Expenditures." Paper presented at the Federal Economic Statistics Advisory Committee Meeting, Bureau of Labor Statistics, Washington, DC, March 21, 2003.
- Government of India, *National Accounts Statistics, Ministry of Statistics and Programme Implementation*, http://mospi.nic.in/national_account_main.htm (accessed March 1, 2003).
- Groves, Robert M., and Mick P. Couper, *Nonresponse in Household Interview Surveys* (New York: Wiley, 1998).
- Kakwani, Nanak, "Lorenz Curve" (vol. 3, pp. 242–244), in J. Eatwell, M. Milgate, and P. Newman (Eds.), *The New Palgrave: A Dictionary of Economics* (London and Basingstoke: Macmillan, 1987).
- Keidel, Albert, "China's GDP Expenditure Accounts," *China Economic Review*, 12: 4 (2001), 355–367.
- Kulshreshtha, A. C., and A. Kar, "Estimates of Food Consumption Expenditure from Household Surveys and National Accounts," World Bank Working paper, <http://www.worldbank.org/indiapoverty-workshop> (2002).
- Kuznets, Simon, "Economic Growth and Income Inequality," *American Economic Review*, 45: 1 (1955), 1–28.
- Maddison, Angus, *Chinese Economic Performance in the Long Run* (Paris: OECD, 1998).
- Minhas, B. S., "Validation of Large-Scale Sample Survey Data: Case of NSS Household Consumption Expenditure," *Sankhyā, Series B*, 50: Supplement (1988), 1–63.
- Minhas, B. S., and S. M. Kansal, "Comparison of NSS and CSO Estimates of Private Consumption: Some Observations Based on 1983 Data," *The Journal of Income and Wealth* 11 (1989), 7–24.
- Mistaen, Johan A., and Martin Ravallion, "Survey Compliance and the Distribution of Income," Policy Research Working Paper 2956. World Bank, Development Research Group, Washington, DC, 2003.
- NSSO Expert Group on Sampling Errors, "Suitability of Different Reference Periods for Measuring Household Consumption: Results of a Pilot Study," *Economic and Political Weekly* (January 25, 2003), 307–321.
- OECD, *Measuring the Non-observed Economy: A Handbook* (Paris: OECD, 2002).
- Office for National Statistics, *Annual Abstract of Statistics* (London: HMSO, 2003).
- Pyatt, Graham, "Development and the Distribution of Living Standards: A Critique of the Evolving Data Base," *Review of Income and Wealth*, 49: 3 (2003), 333–358.
- Ravallion, Martin, "Growth, Inequality, and Poverty: Looking beyond the Averages," *World Development*, 29:11 (2001), 1803–1815.
- "Measuring Aggregate Welfare in Developing Countries: How Well Do National Accounts and Surveys Agree," this REVIEW, 85 (3), 645–652.
- Rawski, Thomas G., "What is Happening to China's GDP Statistics?" *China Economic Review*, 12: 4 (2001), 347–354.

- Sala-i-Martin, Xavier, "The disturbing 'rise' of global income inequality," NBER working paper no. 8904 (2002).
- Srinivasan, T. N., and Pranab Bardhan, *Poverty and Income Distribution in India* (Calcutta: Statistical Publishing Society, 1974).
- Sundaram, K., and Suresh Tendulkar, "NAS-NSS Estimates of Private Consumption for Poverty Estimation: A Further Comparative Estimation," *Economic and Political Weekly* (January 25, 2003), 376–384.
- Taylor, Lance, and Edmar L. Bacha, "The Unequalizing Spiral: A First Growth Model for Belindia," *Quarterly Journal of Economics*, 90: 2 (1976), 197–218.
- Triplet, Jack E., "Measuring Consumption: The Post-1973 Slowdown and the Research Issues," *Federal Reserve Bank of St Louis Review* (May/June 1997), 9–42.
- Wu, Harry X., "China's GDP Level and Growth Performance: Alternative Estimates and the Implications," *Review of Income and Wealth*, 46: 4 (2000), 475–499.

APPENDIX

Log-normally Distributed Income with Selective Compliance

Suppose that x is the logarithm of income, and that x is normally distributed in the population with mean μ and variance σ^2 . The compliance probability as a function of income is given by equation (1) in the main text. In general, if the true density is $f(x)$, the density function of the truncated distribution is given by

$$\hat{f}(x) = \frac{p(x)f(x)}{\int_0^\infty p(s) dF(s)} = \phi(x)f(x). \quad (\text{A-1})$$

In this normal case with the response function given by equation (1) in the main text, the truncated density is

$$\hat{f}(x) = \begin{cases} \frac{1}{\bar{p}\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], & x \leq \mu - \theta\sigma, \\ \frac{1}{\bar{p}\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right] \\ \times \exp[-\alpha(x-\mu-\theta\sigma)], & x \geq \mu - \theta\sigma, \end{cases} \quad (\text{A-2})$$

where \bar{p} is the population-average compliance probability. The second part of equation (A-2) can be rewritten

$$\hat{f}(x) = \frac{1}{\bar{p}\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2}\left(\frac{x-\mu-\alpha\sigma^2}{\sigma}\right)^2\right] \times \exp\left(-\alpha\theta\sigma + \frac{1}{2}\alpha^2\sigma^2\right), \quad x \geq \mu - \theta\sigma. \quad (\text{A-3})$$

If we integrate $\hat{f}(x)$ over the full range of x , we can derive an expression for the mean compliance probability

$$\bar{p} = \Phi(-\theta) + \Phi(\theta - \alpha\sigma) \exp\left(-\alpha\theta\sigma + \frac{1}{2}\alpha^2\sigma^2\right), \quad (\text{A-4})$$

where the first term comes from integrating the first part of equation (A-2), and the second from integrating equation (A-3). These three equations completely characterize the truncated density $\hat{f}(x)$.

If we substitute equation (A-4) into (A-3), we get the density of a normal distribution with mean $\mu - \alpha\sigma^2$ and variance σ^2 , scaled by the factor

$$\frac{\exp\left(-\alpha\theta\sigma + \frac{1}{2}\alpha^2\sigma^2\right)}{\Phi(-\theta) + \Phi(\theta - \alpha\sigma) \exp\left(-\alpha\theta\sigma + \frac{1}{2}\alpha^2\sigma^2\right)} = \frac{1}{\Phi(-\theta) \exp\left(\alpha\theta\sigma - \frac{1}{2}\alpha^2\sigma^2\right) + \Phi(\theta - \alpha\sigma)}. \quad (\text{A-5})$$

As θ becomes large, the second term in the denominator on the right goes to 1, while the first term goes to 0. Hence, for large θ , with little of the density to the left of $\mu - \theta\sigma$, the truncated density is approximately equation (A-3), which is approximately

$$\hat{f}(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2}\left(\frac{x-\mu-\alpha\sigma^2}{\sigma}\right)^2\right], \quad x \geq \mu - \theta\sigma, \quad (\text{A-6})$$

so that the truncated distribution of log income is also normal, with the same variance σ^2 as the true distribution, but with mean $\mu - \alpha\sigma^2$ instead of μ .