CHAPTER 11

The Global Distribution of Income

Sudhir Anand^{*}, Paul Segal[†]

^{*}Department of Economics, University of Oxford, Oxford, UK; and Department of Global Health and Population, Harvard School of Public Health, Boston, MA, USA

[†]King's International Development Institute, King's College London, London, UK

Contents

11.1. Introduction	938
11.2. Why Study the Global Distribution of Income?	939
11.3. Which Global Distribution of Income?	941
11.4. Data	945
11.4.1 Household Surveys and National Accounts	945
11.4.2 Top Income Data	948
11.4.3 PPP Exchange Rates	948
11.4.4 Estimation Errors	952
11.5. Estimating the Global Distribution of Income	953
11.5.1 Global Inequality Estimates With and Without Top Income Data	955
11.5.2 Comparison of Alternative Estimates of Global Inequality	959
11.5.3 Global Inequality Estimates Using NA Means, Without Top Income Data	960
11.6. Between- and Within-Country Inequality	963
11.7. Relative and Absolute Global Inequality	967
11.8. Global Poverty	968
11.8.1 Methodology	968
11.8.2 Poverty Estimates	971
11.9. Conclusion	973
Acknowledgments	975
Appendix. Estimates of Global Inequality Based on the Common Sample over Time	975
References	977

Abstract

This chapter investigates recent advances in our understanding of the global distribution of income, and produces the first estimates of global inequality that take into account data on the incomes of the top one percent within countries. We discuss conceptual and methodological issues – including alternative definitions of the global distribution, the use of household surveys and national accounts data, the use of purchasing power parity exchange rates, and the incorporation of recently available data on top incomes from income tax records. We also review recent attempts to estimate the global distribution of income. Our own estimates combine household survey data with top income data, and we analyze various aspects of this distribution, including its within- and between-country components, and changes in relative versus absolute global inequality. Finally, we examine global poverty, which is identified through the lower end of the global distribution.

Keywords

Global inequality, purchasing power parity exchange rates, household surveys, national accounts, top incomes, global poverty

JEL Classification Codes

D63, E01, I32

11.1. INTRODUCTION

As the world has become increasingly interconnected through trade, investment, migration and communication, people's interest in and knowledge of international comparisons of living standards has grown. Correspondingly, the global distribution of income has become the subject of numerous research papers and articles, and commentaries in the media. In the popular imagination it seems self-evident to be of interest that great wealth and great poverty coexist in the world. In this chapter we examine the concept of global inequality, the normative motivations for studying it, and the available evidence on the global distribution of income. Widely varying estimates of global income inequality have been published, using a variety of data and methodologies. We critically discuss the different approaches and assumptions behind them, with a view to determining what we believe to be best practice. We also construct a global distribution using both household surveys and top income shares from tax data.

Inequality is a broad concept, and the global distribution of income allows various interpretations. For this reason we start by clarifying different conceptions of the global distribution of income. The distribution of primary interest for us, and the subject of most of this chapter, is that among individuals in the world, each assigned his or her per capita household income. This is what we will refer to as the global distribution of income. But other distributions of global income are also of interest for certain questions. Studies of economic growth and convergence, for instance, are based on changes in the distribution among countries of per capita national income, which is a type of global income distribution that is only indirectly related to the global distribution of income among individuals.

Because individuals around the world are naturally partitioned by country of residence, we examine the between-country and within-country components of global inequality, which can have different definitions depending on the inequality measure used. Although we do not discuss the causes of changes in global income inequality, this decomposition provides a breakdown of those changes, allowing us to isolate the contributions of differential growth in per capita income across countries, and of changes in inequality within countries. This decomposition is a necessary precursor to any causal explanation because one would expect different mechanisms to explain the two components.

Studying the global distribution of income raises difficult empirical and measurement issues. To compare real incomes across countries one needs to convert them using purchasing power parity (PPP) exchange rates, rather than market exchange rates, to take

account of aggregate price differences between countries. There are different methods for calculating PPP exchange rates, which have their respective merits and are used by different studies. All methods depend on the price surveys conducted by the International Comparison Program (ICP). In some cases those price surveys have themselves been controversial. We do not discuss PPP exchange rates in detail (q.v. Anand and Segal, 2008), but we highlight some of the features and controversies that are most relevant for studying the global distribution of income.

Another empirical controversy concerns the measurement of mean incomes within countries. Any global distribution of income must rely on national household surveys to estimate inequality within countries. But some studies, instead of using the mean incomes recorded by those surveys, have taken the relative distributions implied by them and “scaled” them to national accounts estimates of per capita GDP or household consumption expenditure. We argue that there is no good reason to scale to GDP, but that the use of household final consumption expenditure (HFCE) from the national accounts, which is available for most countries, may provide a useful robustness check. Using HFCE rather than the mean incomes from household surveys changes both the level and trend of estimated global inequality.

Beyond reviewing the conceptual and measurement issues underlying the study of the global distribution of income, the empirical aim of this chapter is to use the best available data to construct global distributions of income based on alternative plausible assumptions. The main innovation is to supplement data from household surveys with newly available estimates of top income shares derived from tax data in a range of countries. These data constitute a significant advance in our understanding of the distribution of income both within countries and globally because individuals at the top of the income distribution are either not represented or are underrepresented in household surveys. Unsurprisingly, their inclusion leads to substantially higher estimates of global inequality.

The chapter continues as follows. Section 11.2 discusses the motivation for the study of global income inequality. Section 11.3 analyzes the different concepts of the global distribution of income. Section 11.4 discusses methodological issues and describes the available data, including the top income share data. Section 11.5 presents our constructed global distributions of income and the corresponding estimates of global inequality. Section 11.6 decomposes global income inequality into between-country and within-country inequality and discusses their significance and evolution. Section 11.7 examines the distinction between relative and absolute inequality and presents some preliminary estimates of absolute global inequality. Section 11.8 turns to the estimation of global poverty and considers its level, trends, and regional concentration. Section 11.9 concludes the chapter.

11.2. WHY STUDY THE GLOBAL DISTRIBUTION OF INCOME?

Interest in global inequality reaches far beyond academia and has increased dramatically in recent years—among activists and NGOs, the news media, and national and international

institutions and policymakers. This is in part due to the perception that the benefits of rapid economic growth in recent decades, which has coincided with a period of rapid globalization, have been distributed highly unequally. Thus, the worldwide “Occupy” movement launched in 2011, with its slogan “We are the 99%,” has focused on the sharply increasing concentration of income and wealth among the top 1% of income recipients compared to the other 99%. In the news media, *The Economist* has described growing inequality as “one of the biggest social, economic and political challenges of our time” (Beddoes, 2012). At the 2012 World Economic Forum meeting at Davos, “severe income disparity” was featured as the single most likely global risk, and with one of the highest potential impacts.¹ Again at Davos in 2013, Christine Lagarde, managing director of the International Monetary Fund, stated that “[e]xcessive inequality is corrosive to growth; it is corrosive to society. I believe that the economics profession and the policy community have downplayed inequality for too long” (Lagarde, 2013).

There is indeed a positive case for being concerned about the consequences of inequality for economic growth and social cohesion; crime rates and population health, for instance, have been linked to income inequality within countries.² To the extent that such (within-) country inequality contributes to global inequality, there will be a corresponding concern about the latter. One might equally be concerned about the “corrosive” effects of global inequality itself. Davos, where Lagarde made her comments, is a meeting place of the global élite (i.e., those at the top of the global income distribution, and not just their respective national distributions).

The normative case for studying global inequality seems obvious to some, but it is contested by philosophers who believe that the distribution of income among individuals can be a matter of justice only if they share a government. Nevertheless, even these philosophers typically agree that “there is some minimal concern we owe to fellow human beings threatened with starvation or severe malnutrition and early death from easily preventable disease,” and that therefore “the urgent current issue is what can be done in the world economy to reduce extreme global poverty” (Nagel, 2005, p. 118). In itself this warrants study of at least the lower end of the global distribution of income.

An alternative understanding of justice may lead to a normative concern about global inequality. Some cosmopolitan political theorists argue that egalitarian principles apply equally at the global as at the national level simply because all human beings are entitled to equal respect and concern.³ On this view, national borders are not relevant to an ethical concern with inequality.

It can also be argued that the institutional arrangements that exist in the global economy—the international rules and organizations that govern the flows of goods,

¹ World Economic Forum (2012), reported by Tett (2012).

² See, for example, Pickett and Wilkinson (2010).

³ For discussion see Sen (2000) and Bernstein (2011).

capital, and labor between countries—are sufficient to generate a normative concern for inequality among individuals in the world, even if they fall short of a global government. These international arrangements are largely determined by rich countries and tend to benefit citizens in rich countries at the expense of citizens in poor countries. Rich countries may therefore bear some responsibility for global inequality. Sen (2009, p. 409) puts the issue as follows: “The distribution of the benefits of global relations depends not only on domestic policies, but also on a variety of international social arrangements, including trade agreements, patent laws, global health initiatives, international educational provisions, facilities for technological dissemination, ecological and environmental restraint, treatment of accumulated debts (often incurred by irresponsible military rulers of the past), and the restraining of conflicts and local wars.”

In studying the global distribution of income, we need to distinguish between the recognition of inequality and the obligation and capacity to reduce it. Through its domestic policies, a sovereign state can have more influence on national inequality than on global inequality. This might suggest that, from a policy viewpoint, we should assess within-country inequality differently from between-country inequality (see Section 11.6)—especially if international institutions have limited powers to address between-country inequality. In any case, as we improve our understanding of global inequality, we will be in a better position to diagnose its causes and discuss ways of mitigating it.

In this chapter we take the global distribution of income to be of intrinsic interest. We will analyze various aspects of this distribution, including its within- and between-country components, and also its lower end, which is needed to identify global poverty. Constructing the global distribution of income can be the first step in a broad exercise, which should ultimately permit us to examine many different aspects of global inequality—such as the extent to which gender, ethnicity, education and other socioeconomic variables contribute to global inequality, the characteristics of the global poor, and the composition of the global top 1%.

11.3. WHICH GLOBAL DISTRIBUTION OF INCOME?

Our starting point must be to clarify what we mean by the global distribution of income. Following Milanovic (2005) and Anand and Segal (2008), we can define four concepts of the global distribution of income and their associated levels of inequality, distinguished by the population unit and the income concept (which may be a measure of consumption expenditure) to which they refer. The four concepts of global distribution are relevant to addressing quite different questions, as we discuss in this section. We must also decide on the numéraire to make the income concept comparable across countries, and use either market exchange rates or PPP exchange rates. PPP exchange rates account for the fact that one U.S. dollar will typically buy less in the United States than one dollar’s worth of, for example, Indian rupees purchased on the currency markets, will buy in India. Later

we will discuss different approaches to PPP exchange rates and some of the complications that arise in estimating and using them. Which exchange rate is appropriate will depend on the question being asked.

Our first concept of the global distribution of income, denoted *concept 0*, is the distribution of global income by country. In other words, the “population unit” is the *country* and the “income concept” is the (total) *national income* of the country. Thus India and Canada, both with GDPs of US\$1.8 trillion in 2012, count as equal, despite the fact that India has a population of 1237 million and Canada only 35 million. It is this concept 0 global distribution that is most relevant for questions of geopolitics and market access. In international negotiations over trade rules and macroeconomic policies it is a country’s total economic size that tends to determine its bargaining power. For such questions, it is a country’s weight in international markets—its command over internationally traded goods and services, or financial assets—that matters, and hence income at market exchange rates is likely to be relevant. One might refine the measure, of course, depending on the geopolitical question at hand. For example, in matters concerning global energy markets, countries with relatively small economies but large fuel exports tend to be important.

Next, the *concept 1* global distribution again takes the *country* as the population unit, but now the income concept is the *national income per capita* of the country, not its total national income.⁴ This is the concept typically used in analyses of economic growth, and in particular of economic convergence, where the question is how the set of characteristics and policies associated with a given country affects its per capita income growth rate. Because it is real output that is of interest in this case, income levels will be measured at PPP exchange rates.

In the *concept 2* global distribution, the population unit is the *individual*, and the income concept is again *national (household) income per capita*. (This is equivalent to taking the country as the population unit, as in concept 1, but weighting each country by the size of its population.) It is not obvious why this concept of global inequality would be intrinsically interesting, but some older studies have analyzed its evolution over time, mainly because of the ready availability of data on national income or GDP per capita (Boltho and Toniolo, 1999; Firebaugh, 1999, 2003; Melchior et al., 2000). Concept 2 is of instrumental interest, however, through its relationship with concept 3, which is the focus of this chapter.

The *concept 3* distribution also takes the *individual* as the population unit, but the income concept is the *per capita income of the household to which the individual belongs*—under the assumption of equal sharing of household income (or consumption expenditure). It is the global analogue of the type of distribution typically used to calculate inequality within

⁴ It follows that this *concept 1* distribution—unlike the concept 0 distribution—is not a “distribution of (total) global income among countries.”

Table 11.1 Concepts of global income distribution and inequality

	Unit of analysis	Ranking variable	Numéraire
Concept 0	Country	National income	US\$ or PPP\$
Concept 1	Country	National income per capita	PPP\$
Concept 2	Individual	National (household) income per capita	PPP\$
Concept 3	Individual	Household income per capita of individual	PPP\$

countries.⁵ Henceforth we use the terms “global distribution of income” and “global inequality” without further qualification to refer to their *concept 3* counterparts. Because it is real income or consumption that we are interested in, national currencies will be compared using PPP exchange rates. Concept 3 is also the only concept that tells us something directly about global welfare.

Concept 2 global inequality can be seen as the between-country component of concept 3 global inequality. Concept 2 inequality tells us what concept 3 inequality would be if there were no inequality within countries and each person in a country received the national (household) income per capita of that country. For decomposable measures of inequality, concept 3 inequality will then be equal to concept 2 inequality plus a weighted average of inequality within countries (the within-country component of concept 3 inequality). We will discuss these distinctions further when we present our calculations later in this chapter.

Table 11.1 summarizes the four global income distributions defined in terms of unit of analysis (population unit), associated ranking variable (income concept), and numéraire. It is important to emphasize that the four different concepts of global inequality can move in different directions. It should be immediate from the decomposition just mentioned that concepts 2 and 3 can move in different directions: a modest fall in between-country (i.e., concept 2) inequality may coexist with a rise in concept 3 global inequality if within-country inequality increases sufficiently.

Moreover, the same changes in national income may have opposite effects on different concepts of inequality. For example, China is the second-largest economy in the world, both in PPP\$ and in current US\$, but its per capita GDP in 2012 was PPP\$7960, below both the unweighted mean GDP per capita across countries of PPP\$12,300 and the

⁵ In country studies of income inequality, an adjustment for differential needs and economies of scale in household consumption is sometimes made by taking account of the age and sex composition of a household, in addition to its size. This is done through “equivalence scales,” which allow the calculation of the number of “equivalent adults” in a household. Each individual in the household is then assigned the household’s income per equivalent adult. Given the type of survey data at our disposal, it is not possible to estimate the number of equivalent adults in each household and rank individuals by their household income per equivalent adult. Hence, like other studies of interpersonal global inequality, we simply rank individuals by their household income per capita.

population-weighted mean across countries of PPP\$10,260.⁶ The fact that China's above-average total national income has been growing much faster than the world average therefore implies that China is a disequalizing force for concept 0 inequality. However, the fact that its below-average GDP per capita is growing faster than the world means of both unweighted and population-weighted GDP per capita is an equalizing force for concepts 1 and 2 global inequality, respectively. The latter implies that it is also an equalizing force for concept 3 global inequality.

Consider now the notion of "convergence" in the literature on economic growth, which is the closest that many economists get to thinking about global inequality. There are two commonly used definitions of "convergence," namely beta and sigma convergence. Beta convergence means that when a country's growth rate is regressed on its national income per capita, the coefficient on income is negative and significant.⁷ Thus, on average, countries with higher per capita national income (where per capita GDP is the measure typically used in these studies) have lower growth rates. Sigma convergence means that the dispersion across countries of per capita national income declines over time, often measured by the standard deviation of the logarithm of per capita national income. Both therefore refer to the concept 1 global distribution with the country as the population unit and per capita national income as the income concept.

In their survey of growth econometrics, Durlauf et al. (2009, p. 1098) state that sigma convergence has "a natural connection to debates on whether inequality across countries is widening or diminishing." If "inequality across countries" refers to concept 1 inequality, then it is a tautology that sigma convergence will measure "whether inequality across countries is widening or diminishing."⁸ However, sigma convergence or divergence has no necessary connection to any other concept of inequality across countries (e.g., concept 0 inequality as seen in the China example) or to global inequality across a different population unit (e.g., individuals). A rise in the dispersion across countries of per capita national income (i.e., concept 1 inequality) may be associated with a fall in concepts 2 and 3 global inequality, as the following example demonstrates.

The Philippines has a population of 97 million people, and a per capita GDP of PPP\$3800. Its per capita GDP is below both the unweighted and the population-weighted world means of PPP\$12,300 and PPP\$10,260, respectively, noted earlier. There are 35 countries that both have populations below 5 million and also, like the Philippines, have per capita GDP below the respective unweighted and population-weighted world means. For the purposes of sigma convergence, each of these 35 small

⁶ PPP\$ are in 2005 prices, and data are from *World Development Indicators*, <http://data.worldbank.org/data-catalog/world-development-indicators>.

⁷ *Conditional* beta convergence means that the coefficient on income is negative and significant when other variables are controlled for in the regression.

⁸ One concern about this argument is that the standard deviation of log-income is not a good inequality measure, as it does not satisfy the principle of transfers at the top end of the income distribution (Sen, 1973).

countries has the same weight as the Philippines, yet their combined populations amount to 57 million, below that of the Philippines.⁹ Now imagine that there is sigma divergence, where all other countries are growing at a common rate, but the Philippines is growing faster and the 35 small countries are growing slower than this common rate. Global inequality is increasing according to concept 1 because while one country (the Philippines) whose per capita GDP is below the unweighted world mean is converging to the world mean, 35 other countries whose per capita GDP is below the world mean are diverging from it. But global inequality may be decreasing according to concept 2 because the convergence of the Philippines' large population toward the weighted world mean outweighs the divergence of the populations of the 35 small countries away from the weighted world mean. Assuming that inequality within countries is unchanged, global inequality may therefore also be decreasing according to concept 3.

We conclude that global inequality *tout court* is an underspecified concept, and estimates of different definitions of global inequality can move in different directions—as we find in our empirical estimates in Sections 11.5 and 11.6.

11.4. DATA

11.4.1 Household Surveys and National Accounts

Household surveys are the most widely available source of data for estimating income distributions within countries, and it is the great expansion in their global coverage that has permitted estimates of global inequality. One could in principle use census data or other sources—but in practice these are available for far fewer country-years than household surveys. Survey coverage has expanded dramatically in the last 30 years; the surveys used by the World Bank to estimate global poverty in 1981 covered only 51.3% of the population of the developing world, whereas in 2005 they covered 90.6% (Chen and Ravallion, 2008).

Although there is no credible alternative to using household surveys for estimates of global inequality, they do suffer limitations. Beyond the obvious sampling and measurement errors, surveys may suffer from biases due to underreporting of incomes by the rich and undersampling of both very rich and very poor households. Most important for our purposes, differences in definitions and coverage mean that different surveys are typically not strictly comparable with one another (see Anand and Kanbur, 1993, pp. 33–36). Atkinson and Brandolini (2001) described such problems in the Deininger and Squire database, which collates estimates of inequality within countries; Anand and Segal (2008) discussed these issues in the context of measuring global inequality, observing that in some surveys incomes are gross-of-tax and in others net-of-tax; some refer to cash

⁹ GDP and population data are from *World Development Indicators*, <http://data.worldbank.org/data-catalog/world-development-indicators>.

incomes, whereas others include certain items of income-in-kind; some impute the rental value of owner-occupied housing, whereas others do not. Moreover, all global data sets of household surveys combine surveys of income and of consumption expenditure. There is no reliable way to infer an income distribution from an expenditure distribution, or vice versa, so one simply has to live with the noncomparability. For brevity we will refer to “income or consumption expenditure distributions” as “income distributions.”

The World Bank’s Living Standards Measurement Surveys, initiated in 1980, have been instrumental in increasing both the quantity of survey data available and its quality. The Luxembourg Income Study (LIS) specifically attempts to harmonize survey data to ensure their comparability, and the LIS data set currently covers 47 countries. Still, non-comparability cannot be avoided in a global data set of household surveys, which cover most of the world’s population.

Although all recent studies of global inequality use survey data for estimates of within-country inequality, most then “scale” the within-country distributions to national accounts estimates of mean income or consumption expenditure. For instance, Chotikapanich et al. (1997), Dowrick and Akmal (2005), Sala-i-Martin (2006), and Schultz (1998) use the Deininger and Squire (1996) inequality database for estimates of *relative* inequality within countries and peg the relative distributions around an absolute mean from the national accounts.¹⁰ Milanovic (2002, 2005, 2012) and Lakner and Milanovic (2013) are the only studies we know of that estimate global income inequality using levels of income or expenditure directly from surveys, rather than scaling relative distributions to NA means (though Lakner and Milanovic do use NA means in imputing top incomes, as we discuss later). The World Bank also uses absolute incomes from household surveys for its estimates of global poverty (Chen and Ravallion, 2001, 2008, 2012).¹¹ The distinction between using survey data directly and scaling them to national accounts categories matters because both the levels and rates of change of global inequality and poverty can vary substantially (Deaton, 2005).

For studies that use household surveys only for their relative distributions and scale them to national accounts means, there are two widely available estimates of NA “mean income”: per capita GDP and per capita household final consumption expenditure (HFCE). In principle, one would want to use the category of personal income, but countries do not usually report this category. Most studies of global inequality simply use per capita GDP as a proxy for individual mean (per capita household) income.¹²

¹⁰ See Anand and Segal (2008) for detailed descriptions of their methodologies.

¹¹ The World Bank’s own estimates are based on unit record data. These data are available to the public only in coarser, grouped form from the World Bank’s Povcalnet Website: <http://iresearch.worldbank.org/PovcalNet/index.htm>.

¹² These studies are Bourguignon and Morrisson (2002), Bourguignon (2011), Sala-i-Martin (2006), Dowrick and Akmal (2005), Schultz (1998), Chotikapanich et al. (1997), and Korzeniewicz and Moran (1997). Dikhanov and Ward (2002) use “personal consumption expenditure,” which we take to mean HFCE.

As argued in Anand and Segal (2008, pp. 66–68), if it is national household consumption expenditure that one wishes to measure, then there is no reason to use GDP when HFCE is available. Moreover, GDP is also a poor measure of *household income*: GDP includes depreciation, retained earnings of corporations, and the part of government revenue (taxes) that is not distributed back to households as cash transfers. Deaton (2005, p. 4) noted 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.” In the case of the United States, which is one of the few countries that does report measures of aggregate household income (referred to as “personal income”), it amounts to only about 70% of GDP. Deaton estimates that, across 272 surveys of household income from around the world, survey household income amounts on average to only 57% of GDP, but equals 90% (101% population-weighted) of HFCE from National Accounts.

The question remains, however, whether one would want to use *any* National Accounts figures when mean household income (or consumption) is available in the surveys themselves, which are the source of the income (or consumption) distribution for countries. We saw earlier that surveys have their own problems. But they are at least a direct measure of the variable of interest. HFCE, on the other hand, includes the category of “non-profit institutions serving households” (e.g., religious organizations and political parties), and suffers from being calculated as a residual of aggregate consumption minus estimates of firms’ consumption and government consumption. Errors in any of the latter magnitudes will translate into errors in estimates of HFCE.¹³

New evidence on national accounts data in low-income countries casts more general doubts on their reliability. Jerven (2013) noted that Ghana revised its GDP upward by 60.3% in November 2010 owing to a change in base year,¹⁴ and argues that similarly large revisions are to be expected in other sub-Saharan African countries.¹⁵ Young (2012) also found that national accounts provide a poor measure of growth in sub-Saharan Africa and produces independent estimates of consumption growth based on data from Demographic and Health Surveys.¹⁶

Most of our analysis that follows will refer to the global distribution based on mean incomes from household surveys, but we also calculate global inequality where mean per capita household income is taken to be equal to per capita HFCE as reported in the National Accounts and compare the differences in the results.

¹³ See Anand and Segal (2008, p. 68) for detailed discussion.

¹⁴ The 1993 base-year estimates excluded parts of the economy that were important in the new base year of 2006 (Jerven, 2013, p. 11).

¹⁵ These countries are Nigeria, Uganda, Tanzania, Kenya, Malawi, and Zambia. Jerven’s explanation is that many of these countries suffered drastic cuts to statistical services in the 1980s and 1990s.

¹⁶ Note, however, that Young’s method of inferring aggregate consumption from data on assets has been criticized by Harttgen et al. (2013).

11.4.2 Top Income Data

Perhaps the most important recent innovation in estimating national income inequality has been the collation of data on top income shares from income tax records. These estimates present the incomes of the top 0.1%, top 1%, and top 10% as a share of “control” income, where control income is an estimate of total personal income in the economy (not just taxable income). They are important primarily because they make a substantial difference to estimated inequality. Household surveys typically undersample (exclude) the richest individuals or underreport their incomes, or both. In the United States in 2006, for instance, tax data excluding capital gains imply a top percentile share of 18.0%, whereas survey data imply a share of 13.7%. Using data for 2006, the U.S. Gini based on household survey data (the Current Population Survey) is 0.470, whereas correcting the top percentile’s income using the tax data raises it by nearly 0.05 to 0.519. Moreover, the increase in the U.S. Gini from 1976 to 2006 using survey data alone (corrected for a change in definition) was 0.053, which more than doubles to an increase of 0.108 using the top income data (including capital gains; Atkinson et al., 2011, p. 31; see Burkhauser et al., 2009 for further discussion of U.S. data).

Atkinson et al. (2011, pp. 4–5) describe the top income data in detail, and discuss their limitations. These include the fact that the income shares refer to gross income before tax; the data vary with respect to the unit of observation, some referring to individuals and others to households; in some cases they are not consistent over time, as tax regimes change; and they may be biased owing to tax avoidance and tax evasion. Although they are typically much better than surveys at capturing capital income, this varies depending on the extent to which capital income is taxed and hence reported in the tax records (Atkinson et al., 2011, p. 35). Alvaredo and Londoño Vélez’s (2013) study of top incomes in Colombia notes that different definitions of the control income, of which top incomes are expressed as a share, lead to somewhat different estimates. For these reasons international comparisons of these top income shares may suffer from inconsistencies. Nonetheless, we will set aside such concerns and use these data on the presumption that excluding them would cause a large negative bias in estimates of global inequality. Clearly, however, these noncomparabilities do add uncertainty to the estimates.

11.4.3 PPP Exchange Rates

International comparisons of living standards require the use of PPP exchange rates to convert national currencies into a common numéraire.¹⁷ Two standard sets of PPPs are publicly available: those produced by the International Comparison Program (ICP) of the World Bank (World Bank, 2008) and those produced by the Penn World Tables (PWT), which also uses the underlying price survey data collected by the 2005

¹⁷ An early discussion of this issue may be found in Berry et al. (1983).

ICP.¹⁸ PPPs for years before and after the “benchmark” year 2005 are derived from each country’s domestic price indices.

The price surveys undertaken for the 2005 ICP were both more detailed and more representative globally than in previous rounds of the ICP. China had never taken part in an ICP before the 2005 round, and India had not taken part since 1985, but both countries were surveyed in the 2005 ICP. Previous estimates of PPPs were therefore based on imputations. Partly for this reason, the results from the 2005 ICP have in some cases led to dramatic changes in estimated GDP. Both China and India were found to have real GDPs nearly 40% lower than previous estimates¹⁹ because prices were found to be higher than previously estimated. In the case of China, at least some of this downward revision appears to have been due to sampling problems: its price surveys took place in cities and their environs and did not cover rural areas. For this reason Chinese prices are likely to have been overestimated, and its real income underestimated. Following Chen and Ravallion (2010), and like Milanovic (2012), we make an adjustment to account for this (described later). Milanovic (2012) found that the revisions in the 2005 ICP make a substantial difference to estimated global inequality, raising the Gini by 4.4–6.1 percentage points over the period 1988–2002 and Theil T by 12.5–16.4 percentage points. Other studies that use the 2005 PPPs are Lakner and Milanovic (2013) and Bourguignon (2011), and we discuss their findings later.

Starting from the vector of prices in each country provided by the ICP, the World Bank and PWT use different methods to calculate PPPs. World Bank PPPs are based on the Elteto–Köves–Szulc (EKS) method, whereas PWT uses the Geary–Khamis (GK) method (both with a variety of adjustments made in the process of estimation).²⁰ EKS arose from a statistical approach to index numbers (Deaton and Heston, 2010) and is a multilateral generalization of the Fisher index for two countries (for further discussion, see Anand and Segal, 2008, p. 71). However, under certain assumptions EKS applied to incomes yields an index of real living standards, or utility, and for this reason Neary (2004) included it as an example of the “economic” approach to index numbers. Under the economic approach it is assumed that observed quantities arise from the optimizing behavior of some representative agent with a well-defined utility function. Real relative

¹⁸ A new ICP with base year 2011 was released recently in June 2014, as this chapter was already in press.

¹⁹ This was calculated by comparing the countries’ respective incomes relative to U.S. income in 2005 at 1993 PPP\$ and at 2005 PPP\$.

²⁰ See Anand and Segal (2008) for details of these two methods, and for discussion of the “Afriat method” used by Dowrick and Akmal (2005) to measure global inequality. The World Bank’s PPPs use EKS within regions of countries and then link regions using a “ring” of 18 countries with at least two in each region. See Deaton and Heston (2010) for discussion of both the World Bank and PWT PPP methods. These authors also noted that a single global EKS calculation leads to some nontrivial differences compared to the ICP PPPs, including a real GDP in China that is 6.6% higher.

incomes measured using EKS PPPs represent relative utility levels when utility is quadratic (i.e., in these circumstances it is a “true” index).

GK, on the other hand, is an example of the “test” or “axiomatic” approach. The GK index has no interpretation in terms of optimizing behavior, but its putative advantage with respect to EKS is that it passes the test, or obeys the axiom, of *matrix consistency*. That is to say, GK provides a vector of “international prices” for individual goods that enable disaggregation of the economy into subsectors whose values at those prices sum to the total value of the economy. This is not true of EKS, which computes the relative size of aggregate incomes but does not provide a set of international prices with which economies can be consistently disaggregated. If one is interested in analyzing the structure of economies, then matrix consistency would seem to be a useful property. For instance, it is hard to interpret the relative size of manufacturing in two different countries when manufacturing plus nonmanufacturing within each country does not add up to 100% of its economy.

Matrix consistency would seem less relevant, however, when our concern is international comparisons of living standards. In this case, it is the overall value of consumption, not its composition, that concerns us. More important for our purposes is the drawback of the GK method, which is that it suffers from Gershenshron (or substitution) bias. Because consumers tend to substitute away from goods that are relatively expensive and toward goods that are relatively cheap, valuing the output of both country A and country B at country B’s prices will lead to an overestimation of the income of country A relative to that of country B. The relative prices arising from the standard GK method more closely resemble those in rich countries than in poor countries, leading to an overvaluation of the incomes of poor countries relative to rich countries and therefore to an underestimation of inequality between countries. Ackland et al. (2004) found that the GK method overvalues the incomes of poorer countries compared to EKS. They regress log per capita GDP from GK on log per capita GDP from EKS and find the slope to be 0.94 and to be significantly less than 1.0. Deaton and Heston (2010) found that the Gini for concept 2 (between-country) global inequality, with per capita GDP as the income concept, is slightly higher using EKS than GK, at 0.533 as opposed to 0.527.

Almås (2012) also found that PWT PPPs underestimate global inequality when accounting for both substitution bias and differences in the quality of goods across countries. However, her estimates are based on the strong assumption that “there is a stable relationship between the budget share for food and household income; i.e., there is a unique Engel relationship for food in the world” (Almås, 2012, p. 1094). Deaton and Heston (2010, p. 5) pointed out that “there are many places in the world, such as North and South India, where there are large differences in consumption patterns of food in spite of only modest differences in relative prices.”

Neary (2004) presents a method that he denotes “Geary–Allen International Accounts” (GAIA) for constructing PPPs that is “economic” in the sense of being based on the assumption of optimizing behavior and therefore does not suffer from substitution

bias, but that also satisfies a form of matrix consistency. However, the form of matrix consistency satisfied is not the form that GK satisfies; the sectoral quantities that sum to the value of the whole economy are not the actual observed sectoral quantities, but virtual quantities that a reference consumer, whose preferences are estimated from the data, *would* have chosen. So it is also the case in the GAIA method that observed manufacturing plus observed nonmanufacturing within an economy will not, in general, add up to 100% of the economy.

The theoretical advantage of GAIA over EKS is that it is a “true” index (i.e., produces estimates of relative real incomes that are consistent with optimizing behavior) for a wider range of utility functions. But because all such indices make the false assumption of identical tastes in all countries worldwide, this seems a rather limited benefit. EKS, on the other hand, has the advantage of being relatively transparent. Although GAIA requires the estimation of a demand system, the EKS exchange rate for a country is simply the geometric mean of that country’s Fisher price indices relative to every other country and, as already mentioned, has a natural statistical interpretation that is attractive to national income accountants if not to consumer theorists (Deaton and Heston, 2010).

In our calculations that follow, we use the EKS-based World Bank consumption PPPs from the 2005 ICP. Following Chen and Ravallion (2008, 2010) we make the following adjustments. For both India and China, where the survey data are provided separately for rural and urban strata, we deflate urban incomes relative to rural incomes from price indices used for the construction of domestic urban and rural poverty lines. For India we assume that the World Bank estimated PPP is a weighted average of the urban and rural PPPs. For China we assume that the reported PPP is for urban areas and adjust rural prices downward. This is because the price surveys in China in 2005 were restricted to 11 metropolitan areas, which did not include any rural areas (Chen and Ravallion, 2010). The result is a lower overall price level for China, and thus higher average living standards, than those implied by the use of the 2005 ICP.

A limitation to all standard PPP estimates is that they assume all households within a country face the same price level for their expenditure basket. This may be problematic for at least two reasons. First, urban and rural areas typically have different price levels, and although we have taken this into account for China and India, where the urban and rural price surveys are distinct, it is not possible to do so for most countries. Second, different quantiles of a national income distribution will typically consume different baskets of goods and services,²¹ and hence face different costs of living. For instance, the poor may face higher unit costs for a good because they have to buy it in smaller quantities. Moreover, they purchase goods in different proportions from the nonpoor so the prices of goods will have different expenditure weights for them. At the other end of the

²¹ Deaton and Dupriez (2011) discussed this in the case of the poor and have come up with PPPs specifically for estimating global poverty.

distribution, the very rich (such as those captured by the top income data) may tend to buy more goods from outside their country of residence, to which market exchange rates would apply. But to the extent that the very rich spend their income on nontradable goods and services—for example, country estates, urban mansions, and domestic labor within their country of residence—PPPs with different expenditure weights may be more appropriate than market exchange rates.

11.4.4 Estimation Errors

The preceding discussion of the available data indicates that there are several sources of error in estimates of global inequality, including our own. These include sampling errors, which arise from the sample not being representative of the world population. Our global income distribution is constructed as the union of national income distributions, each of which is based on a national household income (or expenditure) survey with a distinct sampling frame and sampling errors (including undersampling of both the rich and the poor in a country). This global distribution is not estimated from a stratified random sample of the world population, so standard methods are not applicable to calculate sampling errors or confidence intervals for estimates of global inequality.

It is important to distinguish sampling errors from other types of estimation error, which arise from imprecise data and invalid or inaccurate assumptions and methods used to calculate global inequality. For example, there are measurement errors in the income or expenditure data in household surveys (e.g., underreporting of incomes of the rich) and in any national accounts data that may be used; there are also estimation errors in the PPP exchange rates used to construct a global income distribution from national distributions. Major revisions in the estimation of PPPs in the 2005 ICP round, discussed earlier, suggest great sensitivity to the assumptions and methods employed. Given such instability, we may expect further revisions in the next set of PPPs from the 2011 round of the ICP.²² Moreover, as mentioned earlier, a single PPP exchange rate for a country may fail to capture differences in price levels faced by households in different quantiles of the income distribution or in different geographical locations in the country.

Bourguignon and Morrisson (2002) estimated global inequality from 1820 to 1992 through the use of inevitably limited data and manifold assumptions. Given the limitations of their data, they simulated “uncertainty” in their mean income (i.e., GDP) numbers and in their country-group distributions (11 data-points for each of 33 countries or groups of countries) and calculated standard errors for global inequality on this basis. Under their simulation assumptions, the resulting standard errors on the global Gini turn out to be small: in 1820 the standard error is 0.9 Gini points, in 1950 it is 0.2 Gini points, and in 1992 it is 0.1 Gini points (where 1 Gini point is 0.01 in the Gini scale of

²² These PPPs were not available at the time of writing.

0.00–1.00). In our view the other sources of error discussed earlier would imply much larger confidence intervals than these standard errors suggest.

11.5. ESTIMATING THE GLOBAL DISTRIBUTION OF INCOME

In this section we present new estimates of the global distribution of income that combine household survey data with top income data. These estimates are constructed from Milanovic's (2012) global distribution data set of household surveys for five "benchmark" years in the period 1988–2005, which we have supplemented with top income estimates from income tax data. Milanovic's data are provided in quantiles, in most cases 20 income groups each comprising 5% of the population. For those countries for which Milanovic (2012, pp. 10–11) has unit record data, he compared inequality based on individual records with that based on the constructed vigintile (5%) shares and found that the underestimation of the Gini using vigintile shares varied from 0.001 to 0.006 with a mean of 0.003. We agree with Milanovic that this is small enough to be inconsequential.

The five benchmark years 1988, 1993, 1998, 2002, and 2005 each have surveys for between 103 and 124 countries and cover between 87% and 92% of the world population and between 95% and 98% of global GDP in PPP\$. The Milanovic data set provides incomes in national currencies, which we convert to our numéraire of international dollars using World Bank PPPs.²³ We thus have incomes from household surveys in PPP\$ for 87% of the world population in 1988, and 90–92% in the later years. There are 67 countries for which we have both survey and PPP data in *all* five benchmark years, which we refer to as the "common sample over time".

As seen in Table 11.2, we have a total of 537 country-years in our data set. Of these, 104 country-years, ranging from 18 to 23 countries in each year, also have income tax data on the share of the top percentile of the population, which we downloaded from the World Top Incomes Database.²⁴ These countries include the three largest developing countries—China, India, and Indonesia; one Latin American country—Argentina; one African country—South Africa; and all the G7 countries.

The rationale for using income tax data for top percentile shares is that household surveys typically fail to capture the incomes of the richest members of society. For example, Székely and Hilgert (1999) found that in most surveys in Latin America the richest

²³ For Soviet republics in 1988 we use Milanovic's calculations based on Milanovic (1998). These are not strictly comparable to World Bank PPPs because they are based on an earlier set of price surveys. For some other countries without PPP exchange rates in the World Bank's World Development Indicators online database, we derived PPPs implicitly from World Bank Povcalnet data, <http://iresearch.worldbank.org/PovcalNet/index.htm2>.

²⁴ The World Top Incomes Database, constructed by Facundo Alvaredo, Tony Atkinson, Thomas Piketty, and Emmanuel Saez. Data downloaded from <http://topincomes.g-mond.parisschoolofeconomics.eu/>.

Table 11.2 Coverage of countries and populations with *both* household surveys and PPP data, 1988–2005

Year	Number of countries	Population in billions (% of world population)
1988	92	4.45 (87)
1993	104	5.06 (91)
1998	109	5.32 (90)
2002	113	5.78 (92)
2005	119	5.95 (92)
Total	537	

Source: Authors' calculations.

individuals had an income no higher than what would be expected of a midlevel manager in an international firm. This suggests that very rich households are simply excluded from surveys, which is the assumption we make in incorporating top income data into our survey distributions. In other words, we assume that the survey data in the Milanovic data set represent only the bottom 99% of the population in each country. Accordingly we multiply the population in each income group in the surveys by 0.99 and append the top percentile with its income share from the tax data (assuming that its share of “control” income is equal to its share of survey income). The exclusion of the top percentile implies that mean income in the surveys is underestimated, and our procedure results in a corresponding increase in mean income for each country.

For those country-years that do not have top income data, we impute top percentile shares on the basis of regression. The income share of the top decile in Milanovic's household survey data is strongly correlated with the income share of the top percentile in the independently estimated top income data. Excluding one visible outlier in the 104 country-years with both Milanovic data and top income data,²⁵ the simple OLS regression coefficient of the income share of the top percentile against the income share of the top decile (on the remaining 103 datapoints) has a t -statistic of 7.46 and an R^2 of 0.36. We then added the original mean income from the surveys as a further regressor. Mean income is found to be highly significant with a t -statistic of 6.69, the top decile share becomes still more significant with a t -statistic of 10.33, and the regression R^2 rises to 0.55.²⁶ We use this latter regression to generate predicted values for the income share of the top percentile for country-years without tax data.

²⁵ The outlier is South Africa for 1993, for which the top decile share in Milanovic's data is exceptionally large at 46%, whereas the top percentile share from income tax data, at 10.3%, is much smaller than would be expected given this.

²⁶ The estimated regression equation is $\text{topone} = -6.8 + 0.51\text{topten} + 0.30\text{meaninc}$ where topone and topten are, respectively, the shares of the top percentile and the top decile in percentage points, and meaninc is mean survey income in PPP\$ thousand. Year dummies and demographic variables were insignificant.

Lakner and Milanovic (2013) take a different approach to imputing top income shares in estimating global inequality between 1988 and 2008.²⁷ Following Banerjee and Piketty's (2010) finding in India that a significant part of the discrepancy between estimates of consumption expenditure in the national accounts and in household surveys can be accounted for by missing or underreported top incomes, Lakner and Milanovic attributed the difference between HFCE and survey incomes (when the former is larger than the latter) entirely to the top decile of the national distribution in each country-year, and add this residual to the income of the top decile reported in the survey. They then calculated a Pareto coefficient for each country-year distribution on the basis of the unadjusted survey income in the ninth decile and the adjusted income in the top decile (following the procedure described in Atkinson, 2007). Assuming this Pareto distribution applies *within* the top decile of each country-year distribution, they estimated income shares for the income groups P90–P95 (i.e., percentile 90 to percentile 95), P95–P99, and P99–P100, yielding 12 income groups per country-year.

An implicit assumption behind Lakner and Milanovic's procedure for imputing top incomes is that HFCE per capita is the *correct* measure of mean consumption expenditure (or income), when it is larger than the corresponding survey mean. We have argued against using national accounts means in Section 11.4 and in Anand and Segal (2008). It should also be noted that Milanovic's (2002, 2005, 2012) own previous estimates of "true" global inequality are based on his assumption that survey means are preferable to national accounts means.

11.5.1 Global Inequality Estimates With and Without Top Income Data

Our results for global inequality are presented in Table 11.3 and Figure 11.1. The first notable finding is the very high level of global inequality. Considering the global distribution with top incomes over the period 1988–2005, the Gini varies between 0.722 and 0.735, MLD (or Theil L) between 1.093 and 1.156, and Theil T between 1.114 and 1.206. The top percentile in the world has a share between 17.3% and 20.7% of global income, and the top decile between 58.5% and 62.0%. The richest percentile in the world have mean incomes almost 21 times the world mean income in 2005, or a mean per capita household income of about PPP\$90,000 in 2005. The threshold for being in the top percentile in 2005 was PPP\$42,000.²⁸

As anticipated, the inclusion of top income data raises the estimated levels of inequality relative to those based on household surveys without top income data. The average

²⁷ Unfortunately, the Lakner and Milanovic data for the benchmark year 2008 were not made available to us for our own calculations of global inequality.

²⁸ For comparison, the threshold defining the top percentile in the United States in 2005 was a total household income of PPP\$342,000, which for a four-person household implies a per capita figure of PPP\$85,500, or approximately double the global threshold.

Table 11.3 Global inequality with and without top incomes, 1988–2005

Year	Income share of top percentile (%)	Income share of top decile (%)	Gini	Between-country Gini (% of global Gini)	MLD	Between-country MLD (% of total)	Within-country MLD (% of total)	Theil T	Between-country Theil T (% of total)	Within-country Theil T (% of total)
With top incomes										
1988	17.3	58.5	0.726	0.649 (89)	1.136	0.886 (78)	0.250 (22)	1.114	0.780 (70)	0.334 (30)
1993	17.6	58.5	0.727	0.636 (88)	1.142	0.836 (73)	0.306 (27)	1.115	0.753 (68)	0.362 (32)
1998	19.0	59.5	0.722	0.632 (88)	1.093	0.780 (71)	0.314 (29)	1.145	0.750 (66)	0.395 (34)
2002	20.6	62.0	0.735	0.649 (88)	1.133	0.830 (73)	0.303 (27)	1.206	0.809 (67)	0.397 (33)
2005	20.7	60.0	0.727	0.633 (87)	1.156	0.806 (70)	0.349 (30)	1.188	0.755 (64)	0.433 (36)
Without top incomes										
1988	11.2	54.8	0.705	0.642 (91)	1.063	0.861 (81)	0.202 (19)	0.967	0.764 (79)	0.202 (21)
1993	11.6	54.9	0.707	0.632 (89)	1.069	0.819 (77)	0.250 (23)	0.976	0.745 (76)	0.231 (24)
1998	13.1	56.9	0.698	0.624 (89)	1.008	0.757 (75)	0.251 (25)	0.969	0.732 (76)	0.236 (24)
2002	14.1	58.5	0.711	0.640 (90)	1.046	0.801 (77)	0.245 (23)	1.027	0.788 (77)	0.239 (23)
2005	14.9	56.5	0.701	0.622 (89)	1.060	0.775 (73)	0.285 (27)	0.977	0.725 (74)	0.252 (26)

Source: Authors' calculations.

Figure 11.1 Global inequality with and without top incomes, 1988–2005. *Note: Estimates with top income data included are denoted “+top”. Source: Table 11.3.*

share of the top percentile over the period 1988–2005 increases from 13.0% on the basis of the surveys alone to 19.0% when top income data are included. Correspondingly, the average top decile share over the period is 56.3% with survey data alone, and 59.7% when top income data are added. Depending on the year, the Gini increases by 3–4%, MLD (or Theil L) by a larger margin of 7–9%, and Theil T by the largest margin of 14–22%. For all measures the increase is greatest in 2005, when the inclusion of top income data raises the global Gini by 4%, MLD by 9%, and Theil T by 22%. These differences in impact reflect the different sensitivities of the measures to income changes at the top end of the distribution.

Turning to changes in inequality with top income data during 1988–2005, the income share of the top percentile rises monotonically from 17.3% to 20.7%. The share of the top decile rises from 58.5% in 1988 to 60.0% in 2005, peaking at 62% in 2002. The Gini coefficient shows very little movement in this period: the highest Gini value is 0.735 (in 2002), which is only 0.013 higher than the lowest Gini value of 0.722 (in 1998), a difference of under 2%. MLD (or Theil L) and Theil T show somewhat larger movements, with the difference between the highest and lowest years for the measures being 6% and 8%, respectively. MLD peaks in 2005 and Theil T in 2002, and for both of these measures inequality rises over the period 1988–2005—for MLD by 1.8% and for Theil T by 6.6%.

The top income data modify both the estimated level of global inequality and its rate of change over time. Although inequality rises over 1988–2005 according to MLD and Theil T with top income data included, inequality is virtually unchanged over the period

according to all three measures when top income data are not included. In the latter case the Gini is 0.705 in 1988 and 0.710 in 2005, MLD is virtually unchanged at 1.06, and Theil T rises marginally from 0.967 to 0.977; however, the income share of the top percentile rises from 11.2% to 14.9%.

The changes in inequality over time are not large compared to changes witnessed in some individual countries. This is particularly so in the case of the Gini coefficient, where the peak-to-trough difference is only 1.3 Gini points with the top income data, compared to a rise, for example, of about 5 Gini points in the United States over the period 1988–2005.²⁹ Moreover, given the different sources of estimation error that we described in Section 11.4, the small changes we find in the global inequality indices may not be statistically significant—particularly in the cases of the Gini and MLD, which are less than 2 percentage points different in 2005 from 1998. However, the rise in the share of the global top percentile, from 17.3% to 20.7% during 1988–2005, seems less trivial; it implies that the incomes of the top percentile increased by 20% relative to mean income—though we note that this is also smaller than the rise in the share of the top percentile in the United States over the same period, from 15.5% to 21.9%.³⁰

In the Appendix we provide analogous results for the “common sample over time” of 67 countries. Whereas the full sample shown earlier comprises between 87% and 92% of the world population depending on benchmark year, the common sample over time comprises between 79% and 82%. As can be seen in Appendix Table 11.A2, the global inequality estimates are very similar to those in Table 11.3 shown earlier. The Gini coefficient for the common sample is never more than 1 percentage point different from that for the full sample, whereas MLD and Theil T are never more than 3 percentage points different. Note that the common sample over time is not necessarily more representative of the global income distribution than our full sample in each benchmark year, and estimates of the level or rate of change of global inequality based on the common sample are not necessarily more accurate.

Our calculations with top income data assume that household surveys do not capture the top percentile of the national income distribution. An alternative way to include the top income data is to assume that surveys are indeed representative of all households, but that they underreport the incomes of the top percentile in the national distribution. This is the assumption made by Alvaredo and Londoño Vélez (2013) and requires a different calculation. Rather than multiply the population of each income group in the survey data by 0.99 and then append the top percentile with its income share from the tax data, on the alternative assumption one simply *replaces* the income of the top percentile in the survey data with that from the tax data. We have performed this calculation as well, and it leads

²⁹ See Atkinson et al. (2011, p. 33, Figure 6), series adjusted with tax data including capital gains.

³⁰ This refers to the share of the top percentile including capital gains, downloaded from the World Top Incomes Database, <http://topincomes.g-mond.parisschoolofeconomics.eu/>.

to marginally lower estimates of global inequality: in the five benchmark years the global Gini is upto 0.4% smaller, and MLD and Theil T are upto 1.2% smaller. However, for the latter two decomposable indices, the within-country component is noticeably smaller, by 3.6–5.2% for MLD and by 2.4–4.1% for Theil T—but the between-country component is about 0.5% larger for both indices.

11.5.2 Comparison of Alternative Estimates of Global Inequality

We saw earlier that only three previous studies have used 2005 PPPs to estimate global inequality: Bourguignon (2011), Milanovic (2012), and Lakner and Milanovic (2013). Milanovic's (2012) estimates of global inequality are directly comparable to our estimates without top incomes in Table 11.3, as they are based on the same survey data and methodology. The only substantial difference we know of is in the PPPs used for countries for which the World Bank does not have data (see footnote 23 for the sources that we use in these cases). Milanovic found the Gini coefficient to vary between 0.684 and 0.707 in the period 1988–2005, whereas in our estimates given earlier it varies between 0.698 and 0.711. However, whereas we find Theil T at virtually the same level in 1988 as in 2005, he found it to rise from 0.875 to 0.982 over the same period.

Lakner and Milanovic (2013), like us, estimated global inequality both with and without imputed top incomes. Their estimates without top incomes also follow the same methodology as Milanovic (2012) and are based directly on survey data. Lakner and Milanovic's estimates of the global Gini without top incomes are close to ours, varying between 0.705 and 0.722 in the period 1988–2008. Their Theil T is slightly higher than ours, varying between 1.003 and 1.049 in the period. Significantly, their MLD shows a marked decline, from 1.142 in 1988 to 1.027 in 2008.

Lakner and Milanovic—like us—found that imputing top incomes leads to higher estimates of global inequality. Their HFCE-based method of imputing top incomes, discussed earlier, raises the global Gini by 3.8–6.3 Gini points, with the difference rising over time in the period 1988–2008.³¹ Nonetheless, their Gini ends the period at almost exactly the same level as it began, declining marginally from 0.763 in 1988 to 0.759 in 2008. This is a much larger effect than we find from adding top income data to the survey data. As we saw in Table 11.3, our method leads to the Gini being approximately 2 Gini points higher in each year. Lakner and Milanovic themselves pointed out that their imputation assumption is rather extreme in some cases. For example, in 2008 in India—the country that appears to have motivated their procedure—they find the survey mean to be only 53% of HFCE per capita, so they attribute the remaining 47% of total HFCE *entirely* to the top decile, adding it to the income of the top decile reported in the survey. This adjustment seems implausibly large to us. Conversely, for China in both 1988 and 2008,

³¹ Lakner and Milanovic do not give estimates of other inequality measures for their distribution with imputed top incomes.

HFCE is smaller than survey income, so no adjustment is made by the authors for under-reporting or undersampling of top incomes.

The final study that uses 2005 PPPs to estimate global inequality is Bourguignon (2011), which—unlike the other studies mentioned in this section—scales within-country distributions to GDP per capita. Bourguignon found the Gini coefficient to decline from approximately 0.70 to 0.66 between 1989 and 2006 (these numbers were read off his Figure 1). This is a substantial decline compared with the findings reviewed earlier of virtually no change in the Gini without top incomes. The main difference between Bourguignon's estimates and the other estimates without top incomes discussed here is that Bourguignon scales to national accounts data. In Section 11.4 we argued that if one uses national accounts data then HFCE is preferable to GDP as an approximation to household income, so we compare estimates based on survey means with estimates based on HFCE means in the next section.

11.5.3 Global Inequality Estimates Using NA Means, Without Top Income Data

In this section we report global inequality estimated by scaling household survey incomes so that the scaled mean is equated to per capita HFCE from NA in each country (in contrast to using incomes directly from the surveys). HFCE figures in PPP\$ are not available for all the country-years for which we have household survey data. In each year, we distinguish between the “full sample” defined as the set of all countries with survey data, and the “common sample” *across data sets* defined as the subset of the full sample countries that *also* have HFCE data in PPP\$ (note that this is *different* from the “common sample over time,” defined earlier). In 1988 the common and full samples are quite different: in that year the countries that have *both* survey data and HFCE data in PPP\$ comprise only 77% of the world population, compared with 87% of the world population for countries in the full sample (see Table 11.2).³² In the other years covered in Table 11.2, the common sample has 3–4% less of the world population than the full sample.³³

For each of our indices, Table 11.4 presents three different global inequality estimates without top income data: first, the full sample estimates as in Table 11.3; second, estimates based on survey data restricted to the common sample; and third, the common sample estimates based on per capita HFCE (as described earlier). We will refer to the first as the

³² For some countries where World Development Indicators (WDI) does not have HFCE data in PPP\$, HFCE is nevertheless available in local currency units (LCUs). By definition, these countries do not have PPP exchange rates in WDI, but for 11 of them we have PPPs (see footnote 23), which we use with the survey data. However, in almost all cases using these PPP conversion rates gives implausible results for HFCE. For this reason we do not use these data.

³³ For the years 1993, 1998, 2002, and 2005 the common sample covers 87%, 87%, 89%, and 89%, respectively, of the world population compared to the full sample percentages in Table 11.2.

Table 11.4 Global inequality estimates using survey means and HFCE means, without top incomes, 1988–2005

Year	Gini			MLD			Theil T		
	Survey means, full sample	Survey means, common sample	HFCE means, common sample	Survey means, full sample	Survey means, common sample	HFCE means, common sample	Survey means, full sample	Survey means, common sample	HFCE means, common sample
1988	0.705	0.721	0.739	1.063	1.148	1.243	0.967	1.024	1.094
1993	0.707	0.706	0.721	1.069	1.074	1.133	0.976	0.970	1.036
1998	0.698	0.699	0.711	1.008	1.010	1.062	0.969	0.973	1.017
2002	0.711	0.710	0.706	1.046	1.039	1.032	1.027	1.022	1.016
2005	0.701	0.698	0.698	1.060	1.044	1.024	0.977	0.967	0.984

Note: Full sample consists of all countries for which household survey data are available, as in Tables 11.2 and 11.3. Common sample in a given year consists of the subset of countries with *both* survey data and HFCE data in PPP\$ in that year.

Source: Authors' calculations.

full sample survey-means estimates, the second as the common sample survey-means estimates, and the third as the common sample HFCE-means estimates. This highlights the fact that in the first two, mean incomes in each country are obtained directly from the surveys, whereas in the third the mean is externally imposed from HFCE data. These estimates exclude the top income data so that we can focus on the differences in global inequality using survey and NA mean incomes. Figure 11.2 plots the Gini coefficients for the three different sets of data.

The most notable difference between the full sample survey-means estimates and the common sample HFCE-means estimates is that whereas the former appear relatively flat, the latter have a clear downward trend. The two estimates are approximately the same in 2002 and 2005, but because of their different starting points in 1988 the full sample survey-means Gini declines by only 0.004 between 1988 and 2005, from 0.705 to 0.701, whereas the common sample HFCE-means Gini declines by 0.041, from 0.739 to 0.698. However, the common sample survey-means estimates indicate that about half this difference is explained by the difference between the full and common sample: the common sample survey-means Gini declines by 0.023, from 0.721 to 0.698.

The second factor that appears to explain the difference in trend for the HFCE-means Gini and the survey-means Gini is the divergent trend for India specifically in comparing survey and HFCE means, a phenomenon that has been examined in detail by Deaton and Kozel (2005). In our data the average annual growth rate of per capita household consumption expenditure in India from 1988 to 2005 was 2.8% according to surveys, and more than twice that at 5.8% according to HFCE from the National Accounts. When we both restrict the estimates to the common sample *and* exclude India from it, the survey-means Gini declines by 0.029, whereas the HFCE-means Gini decreases by a

Figure 11.2 Global Gini without top incomes, using survey means and HFCE means, 1988–2005.
Note: Common sample consists of surveys restricted to country-years with HFCE data. Source: Table 11.4.

similar magnitude of 0.034.³⁴ Thus ensuring a common sample and excluding India virtually eliminates the difference in trend between HFCE-means and survey-means estimates of the Gini.³⁵

Thus the divergence between HFCE-means and survey-means estimates seems to be due to the loss of as much as 10% of the world population in the “common sample” that has *both* survey and HFCE data, and the divergent trends in India. Given this, in our view the decline in global inequality implied by the HFCE calculations is likely to be illusory. We have not examined estimates based on GDP means, as opposed to HFCE means, for the reasons mentioned earlier. Still, these findings do not seem to corroborate Bourguignon’s (2011) result for the period 1989–2006, based on GDP data and discussed earlier, that “inequality decreases, and it decreases at a very fast pace.”

11.6. BETWEEN- AND WITHIN-COUNTRY INEQUALITY

Table 11.3 also presents estimates of between-country and within-country inequality, and Figure 11.3 plots these estimates. Between-country inequality is defined as global inequality under the hypothetical assumption that every individual is assigned his or her country’s mean per capita household income. It suppresses inequality within countries and measures inequality in the global distribution among world citizens where the only source of variation is mean per capita income across countries (in other words, between-country inequality is just concept 2 global inequality). Between-country inequality is well-defined for any inequality index, and we report it in Table 11.3 for the Gini, MLD (or Theil L), and Theil T measures. For the decomposable measures MLD and Theil T, the difference between overall global inequality and between-country inequality is a weighted average of inequality in each country, and is denoted as *within-country inequality*. In the case of MLD (i.e., Theil L), within-country inequality is a population-share weighted average of the MLD in each country, whereas for Theil T it is an income-share weighted average of the Theil T in each country (Anand, 1983, pp. 86–92).

In the case of MLD (Theil L) only, the within-country component has an additional interpretation: it is equal to what global inequality would be under the hypothetical assumption that mean per capita incomes are equalized between countries, while relative inequality is kept constant within each country. In this sense it is a natural complement to the definition of between-country inequality, and for this reason we consider MLD to be strictly decomposable, but Theil T to be only weakly decomposable (Anand, 1983, pp. 198–202).

³⁴ The levels are different, however, with the survey-means Gini declining from 0.701 to 0.671 and the HFCE-means Gini declining from 0.722 to 0.689.

³⁵ This procedure also substantially reduces the difference in trend for MLD and Theil T, but does not eliminate it as effectively as with the Gini.

Figure 11.3 Between-country and within-country global inequality with top incomes, 1988–2005.
Source: Table 11.3.

Considering estimates of global inequality with top incomes, we make four observations from Table 11.3 (top panel). First, between-country inequality is larger than within-country inequality for both the decomposable indices. Between-country inequality ranges between 70% and 78% of overall global inequality for MLD and between 64% and 70% for Theil T.

Second, the inclusion of top income data increases the within-country component substantially, as would be expected. For MLD the within-country component rises by between 23% and 25%, depending on the year, whereas for Theil T it rises by between 57% and 72%. The between-country component also changes because our imputation of the income share of the top percentile increases country mean incomes by different proportionate amounts.

Third, from 1988 to 2005 between-country inequality declines by all three measures, as shown in Table 11.3. For the estimates with top income data, the between-country

Gini falls by 2% from 0.649 to 0.633, the between-country MLD declines 9% from 0.886 to 0.806, whereas the between-country Theil T declines 3% from 0.780 to 0.755.

Fourth, over the period 1988–2005 within-country inequality clearly increases for both decomposable indices as seen in Figure 11.3. For estimates with top income data, the within-country MLD rises by 40% from 0.250 to 0.349, and the within-country Theil T rises by 30% from 0.334 to 0.433.³⁶

The Gini coefficient is not a decomposable measure in either the weak or strong sense. Although we can define the between-country Gini straightforwardly, the residual from overall global inequality cannot be interpreted as within-country inequality (see Anand, 1983, pp. 311–326). However, as with any inequality index, we can answer the question of what happens to the global Gini and to Theil T when country mean incomes are equalized but relative inequality is kept constant within each country (Anand, 1983, p. 201). This question is relevant in assessing the following claims.

On the basis of the fact that between-country inequality is greater than within-country inequality, Sala-i-Martin (2002, p. 39) stated that “the best strategy to reduce world income inequalities is to induce aggregate economic growth in poor countries.” Similarly, Rodrik (2013, p. 12) noted that “the more rapid growth of poor countries since the 1990s is the key behind the recent decline in global inequality,” concluding from this that “aggregate economic growth in the poorest countries is the most powerful vehicle for reducing global inequality.” For economic growth in poor countries to reduce global inequality, it would of course have to be more rapid than growth in richer countries. In this case, the greatest reduction in global inequality that could possibly be achieved without addressing within-country inequality is calculated by eliminating between-country income differences while keeping inequality within each country constant. Conducting this exercise for 2005 with top income data, the Gini would decline from 0.727 to 0.437 and Theil T from 1.188 to 0.433; in the case of the strictly decomposable MLD, the decline is from 1.156 to its within-country component of 0.349. This is a large decline, but global inequality would still remain at about the level of a high-inequality country such as China, where in 2005 we find the Gini to be 0.430, MLD to be 0.367, and Theil T to be 0.324.

In Section 11.3 we pointed out that the concept of sigma convergence in the growth literature has little relationship to any other concept of global inequality. Bourguignon et al. (2004) used GNI per capita and found that what we call concept 2 inequality fell between 1980 and 2002, while concept 1 inequality rose. Similarly, in our data, between-country inequality (i.e., concept 2 global inequality) declines by all measures during 1988–2005, whereas we find “sigma divergence” when we calculate concept 1 global

³⁶ It should be noted that changes over time in population shares of countries in the case of MLD, and income shares of countries in the case of Theil T, will lead to changes in “within-country inequality” even holding inequality constant *within* each country.

Table 11.5 Concept 1 inequality, calculated using per capita incomes from survey data with top incomes

Year	Gini	MLD	Theil T	Sigma (std. dev. of log-income)	World mean per capita income (PPP\$)	China mean per capita income (PPP\$)
1988	0.501	0.538	0.414	1.15	3424	342
1993	0.535	0.574	0.480	1.12	3683	526
1998	0.552	0.594	0.523	1.11	3923	863
2002	0.575	0.655	0.573	1.15	4148	1042
2005	0.578	0.665	0.580	1.17	4364	1916

Source: Authors' calculations.

Figure 11.4 Global inequality without China, based on survey data with top incomes. Source: Authors' calculations.

inequality, as shown in Table 11.5. The three inequality measures increase when applied to the concept 1 distribution: the concept 1 Gini increases from 0.501 to 0.578, MLD from 0.538 to 0.665, and Theil T from 0.414 to 0.580. The standard deviation of (unweighted) log mean income also rises from 1.15 to 1.17.³⁷

Table 11.5 also presents mean per capita survey incomes for the world, and for China separately. Several papers have estimated global inequality excluding China (e.g., Milanovic, 2012; Sala-i-Martin, 2006; and Schultz, 1998), and we present our estimates in Figure 11.4, which include the top income data. They indicate that global inequality without China increases by all three measures: the Gini rises by 0.050, MLD by 0.217 and Theil T by 0.250. We would note, however, that although these estimates are

³⁷ These figures use per capita incomes calculated from survey data with top incomes. Without the top income data, the surveys alone imply much the same trend for concept 1 inequality: the Gini rises from 0.503 to 0.576, MLD from 0.543 to 0.663, Theil T from 0.417 to 0.576, and sigma again from 1.15 to 1.17.

instructive from the point of view of accounting for global inequality and its evolution, they have no global welfare implications because they exclude approximately one-fifth of the world population.

11.7. RELATIVE AND ABSOLUTE GLOBAL INEQUALITY

In an article entitled “What Are We Trying to Measure?” the development economist Dudley Seers (1972, pp. 34–35, endnote 2) made a telling distinction between relative and absolute inequality when he wrote: “Suppose, for example, that a perspective plan specified that [the] per capita income of Brazil doubled in the next thirty years, but assumed no change in distribution or in the proportion unemployed. Then at the turn of the century, a big landowner in the Matto Grosso could run four cars, instead of two, and a peasant in the North-East could eat two kilogrammes of meat a year instead of one. His son might well be still out of work. Could we really call that ‘development’?”. Although relative inequalities in this example have remained unchanged, the absolute differences have grown in proportion to the expansion of the economy. In addition to considering relative global inequality, there is clearly a case for examining absolute global inequality—as noted, for example, by Ravallion (2004).

The first thorough investigation of measures of absolute global inequality is by Atkinson and Brandolini (2010). They posited a “world social welfare function,” which exhibits a changing social marginal valuation of income at different points along the global income distribution (see also Anand and Sen, 2000). The absolute cost of inequality is then expressed in terms of Atkinson’s (1970) concept of “equally distributed equivalent income” for this social welfare function (see also Kolm, 1969). For any income distribution, Atkinson defines the equally distributed equivalent income as that level of income per head, which, if equally distributed, would yield the same level of social welfare as the existing distribution. Then the absolute cost of inequality is the income per head that is “wasted” as a result of inequality (i.e., it is mean income minus the equally distributed equivalent income). (The relative cost of inequality is the absolute cost divided by the mean, which is the definition of Atkinson’s index of relative income inequality.)

For the Gini welfare function, the absolute cost of inequality is mean income μ multiplied by the Gini coefficient G , and the relative cost is simply G (Anand, 1983; Sen, 1973). In general, mean income μ times a relative inequality measure produces the corresponding absolute inequality measure. For the relative global inequality measures G , MLD , and Theil T , we also estimate the absolute global inequality measures μG , μMLD , and μT , respectively, where μ is the world mean income. The world mean income at 2005 PPP\$ is shown in Table 11.6 for the years 1988–2005.

Table 11.7 shows the evolution of absolute global inequality with top income data between 1988 and 2005 as measured by the Gini, MLD , and Theil T , expressed as a ratio of 2005 world mean income calculated from surveys with top incomes added (PPP\$4364).

Table 11.6 World mean income at 2005 PPP\$, alternative estimates

Year	World GNI per capita	World mean income from surveys, with top incomes	World mean income from surveys, without top incomes
1988	6433	3424	3115
1993	6524	3683	3342
1998	7262	3923	3514
2002	7934	4148	3713
2005	8772	4364	3865

Source: GNI per capita from *World Development Indicators*, <http://data.worldbank.org/data-catalog/world-development-indicators>. Estimates of world mean income from surveys, with and without top incomes, are authors' calculations.

Table 11.7 Absolute global inequality using survey data with top incomes added, as a *ratio* of 2005 world mean income (PPP\$4364)

Year	Absolute Gini	Absolute MLD	Absolute Theil T
1988	0.569	0.891	0.874
1993	0.614	0.964	0.941
1998	0.649	0.983	1.029
2002	0.698	1.077	1.146
2005	0.727	1.156	1.188

Source: Authors' calculations.

Over the 17-year period 1988–2005, there has been an unambiguous rise in absolute global inequality according to all three measures. This is unsurprising given the rise in world mean incomes over this period. To prevent a rise in absolute inequality, relative inequality has to decrease at a faster rate than the rise in mean incomes—which seems an unlikely prospect for the global economy.

In Section 11.2 of this chapter we noted and discussed the widespread concern about global income inequality—in terms of both its level and change. Given that there appears to be little movement in relative global inequality (see Section 11.5), whereas there is a significant widening of absolute global inequality (Table 11.7), the widespread concern about inequality may be based on people making comparisons of living standards in *absolute* rather than relative terms.

11.8. GLOBAL POVERTY

11.8.1 Methodology

Like absolute global inequality, global poverty is a measure based on absolute living standards. To measure global poverty, an absolute poverty line is applied to the global distribution of income and the number of individuals below it calculated. This procedure is employed to monitor global poverty over time—including for the first Millennium Development Goal. Chapter 9 of this volume discusses poverty in developing countries

and regions using different poverty lines. The most widely quoted estimates of global poverty for an absolute poverty line are those produced by Chen and Ravallion (2008, 2012) at the World Bank. The Millennium Development Goal refers to consumption poverty, but the limited availability of surveys around the world and over time necessitates that Chen and Ravallion use a mixture of (consumption) expenditure and income surveys.

Chen and Ravallion's (2008, 2012) estimates for the World Bank use the 2005 ICP PPPs, which, as discussed earlier, are preferable to PPPs based on the earlier ICP rounds. Theirs are the only estimates of global poverty based on unit record data from surveys, which are not publicly released, and which are clearly preferable to the grouped data that are available to other researchers. The World Bank's methodology has been criticized for not scaling the survey data to national accounts (NA) means.³⁸

In Section 11.4 on data we discussed the question of whether to use NA means or survey means in the context of estimating global inequality. The methodology used by the World Bank to measure global poverty, like that used by us to estimate global inequality, uses survey data directly to estimate income (or consumption) levels—converted into international dollars using consumption PPPs. As in the case of global inequality, some authors have calculated global poverty by using survey data for within-country relative distributions and NA data for country mean incomes.³⁹ In the context of global inequality we argued that using means directly from surveys is preferable to scaling them to NA levels, and those arguments apply even more for measuring global poverty.

There is a further consideration that makes scaling to NA categories even less appropriate for estimating global poverty. We know that surveys tend to exclude very rich households and/or underreport their incomes, and for this reason they are likely to underestimate mean income or consumption. But this implies that scaling up the income (consumption) of *every* household to ensure that the survey mean is made equal to the NA mean will imply overestimating the income (consumption) of *all* but the richest households. Put another way, the “missing” income of the rich will be inappropriately divided among the entire population. Poverty will therefore be underestimated (for further discussion, see Anand et al., 2010, pp. 13–14).

Turning to the choice of poverty line, the World Bank uses what is commonly known as the “\$1-a-day” line. This was originally defined in World Bank (1990) as PPP\$1-per-day at 1985 PPP. This poverty line was chosen informally as being representative of the poverty lines of the poorest countries, converted into 1985 PPP\$.

The difficulty arises over how to update a 1985-based PPP\$ value.⁴⁰ Within a single country, one would usually update a poverty line by using a price index based on

³⁸ For example, see Sala-i-Martin (2006).

³⁹ See Sala-i-Martin (2006).

⁴⁰ This discussion draws on Anand et al. (2010).

measured inflation. Because the international PPP dollar is indexed to the US dollar in the ICP benchmark year, one might think that all we need to do is to deflate by the US inflation rate. But updating a poverty line that is denominated in PPP\$ is not so simple. As discussed earlier, calculating a set of PPP exchange rates involves the prices of all countries, so that changes in a country's PPP exchange rate will depend on price changes in all countries. Bangladesh's 1985 poverty line in 1985 PPP\$, adjusted for US inflation during 1985–1993, would not be expected to be equal to Bangladesh's 1993 poverty line in 1993 PPP\$.⁴¹

The World Bank updated the 1985 global poverty line to PPP\$1.08 a day at 1993 PPP and now uses PPP\$1.25-a-day at 2005 PPP. This represents a lower rate of inflation than in the US, but as Chen and Ravallion (2001, p. 288) pointed out: “the fact that \$1.08 in 1993 has a US purchasing power less than \$1 in 1985 does not mean that the real value of the poverty line has fallen. Indeed, if we had simply adjusted the \$1 per day line for inflation in the US between 1985 and 1993 we would have obtained a poverty line which is well above the median of the ten lowest poverty lines at 1993 PPP.” The World Bank has chosen consistency with those domestic poverty lines as the most important criterion in setting a global poverty line. This can be justified by arguing that domestic poverty lines will have maintained their real value within their respective countries better than a PPP inflation-adjusted measure. Therefore each time they updated the poverty line, it was derived as the median of the lowest 10 poverty lines in their data set converted into PPP\$ from the most recent ICP (Chen and Ravallion, 2001, 2008). Although this has an obvious logic to it, Deaton (2010) pointed out that the fact that the composition of the bottom 10 countries will change over time can lead to inconsistency: India exited the bottom 10 countries in their data set in the 2005 update owing to its relatively high growth rate, and because its poverty line was relatively low for its income level, this exit led to a rise in the poverty line relative to where it would be with India in the bottom 10 countries. This has the paradoxical implication that a rise in India's income can lead to a rise in estimated global poverty.

A more fundamental challenge to the PPP\$1-a-day poverty line has been made by Reddy and Pogge (2010). They objected to the money metric approach to global poverty measurement, noting that the PPP\$1-a-day poverty line does not correspond to any “achievement concept” or set of capabilities that are common across countries. That is, there is no reason to think that PPP\$1-a-day in one country will enable the same set of achievements—for example, in terms of nutrition or shelter—as PPP\$1-a-day in another country. Although domestic poverty lines are often set according to some achievement concept, this interpretation is lost when a global poverty line is constructed

⁴¹ More generally, national income at PPP\$ calculated in year $t+n$ is not equal to national income at PPP\$ calculated in year t multiplied by intervening domestic growth and deflated by intervening U.S. inflation (see Anand et al., 2010, p. 6).

using standard PPP exchange rates. Reddy and Pogge argued that an explicit achievement-based threshold should be used to define a global poverty line. This would require costing a minimal basic set of capabilities in each country to yield a money-metric poverty line denominated in local currency. Thus the global capability-based poverty threshold would be represented in income space by the set of these national poverty lines, one for each country.⁴² Although this has theoretical attractions, it has not been implemented in practice.

11.8.2 Poverty Estimates

The question of updating the global poverty line remains contentious, and the World Bank reports poverty headcounts for several different poverty lines. Its poverty measurement website Povcalnet⁴³ also allows the user to choose a poverty line for which it then provides estimates. The latest official publication of global poverty numbers (Chen and Ravallion, 2012) presents poverty headcounts using the following poverty lines all at 2005 PPP\$: PPP\$1-a-day, which they describe as “close to India’s (old) national poverty line” and “an exceptionally frugal line even by the standards of the world’s poorest countries” (Chen and Ravallion, 2012, p. 1); PPP\$1.25-a-day, the line derived from domestic poverty lines in poor countries as described earlier; and PPP\$2-a-day. We report the World Bank estimates in Tables 11.8 and 11.9.

Table 11.8 indicates that the first Millennium Development Goal, which was to halve the percentage of people living below “PPP\$1-a-day” (i.e., PPP\$1.25 at 2005 PPP) from its 1990 level by 2015, was almost achieved by 2008: this percentage declined from 43.1% of the developing world to 22.4%. Chen and Ravallion (2012) reported that the goal was in fact achieved in 2010, though the data for 2010 are not fully representative. If the poverty line of PPP\$1-a-day at 2005 prices is used, then the goal was fully achieved in 2008, with the percentage of the poor in developing countries falling from 30.8% in 1990 to 14.0% in 2008 (Table 11.8).

It is clear from Table 11.9 that the distribution of this decline in poverty was highly uneven. The poverty rate in sub-Saharan Africa at the PPP\$1.25 line was only slightly lower in 2008 than in 1981, at 47.5% rather than 51.5%. It is little comfort that the 2008 figure for this region is a larger drop from the 1993 peak of 59.4%. As is well known, a large share of the decline in global poverty is due to China, which managed to cut poverty from 84% in 1981 to 60.2% in 1990 and to 13.1% in 2008. The world excluding China

⁴² One could then use the relative cost of this capability set in different countries to infer the implied “PPP” exchange rates. There is no reason to think that such exchange rates would be similar to extant PPP exchange rates.

⁴³ <http://iresearch.worldbank.org/PovcalNet/index.htm?2>.

Table 11.8 World bank global poverty estimates, 1981–2008

Poverty line	PPP\$1		PPP\$1.25			PPP\$2	
	Number (millions)	Percent of developing world	Number (millions)	Percent of developing world	Number excl China (millions)	Number (millions)	Percent of developing world
1981	1545.3	41.6	1937.8	52.2	1102.8	2585.3	69.6
1984	1369.3	34.7	1857.7	47.1	1137.8	2680.0	68.0
1987	1258.9	30.1	1768.2	42.3	1182.5	2710.2	64.8
1990	1364.7	30.8	1908.6	43.1	1225.5	2864.1	64.6
1993	1338.1	28.7	1910.3	40.9	1277.6	2941.5	63.1
1996	1150.0	23.5	1704.0	34.8	1261.2	2864.8	58.6
1999	1181.9	23.1	1743.4	34.1	1297.0	2937.9	57.4
2002	1096.5	20.6	1639.3	30.8	1276.2	2848.4	53.5
2005	886.1	16.0	1389.6	25.1	1177.7	2595.8	46.9
2008	805.9	14.0	1289.0	22.4	1116.0	2471.4	43.0

Note: All three poverty lines of PPP\$1, PPP\$1.25, and PPP\$2 are at 2005 PPP.

Source: Chen and Ravallion (2012).

Table 11.9 Headcount index of poverty (%) by region, 1981–2008, for poverty line of PPP\$1.25 at 2005 PPP

Region	1981	1984	1987	1990	1993	1996	1999	2002	2005	2008
East Asia and the Pacific	77.2	65.0	54.1	56.2	50.7	35.9	35.6	27.6	17.1	14.3
China	84.0	69.4	54.0	60.2	53.7	36.4	35.6	28.4	16.3	13.1
Eastern Europe and Central Asia	1.9	1.6	1.5	1.9	2.9	3.9	3.8	2.3	1.3	0.5
Latin America	11.9	13.6	12.0	12.2	11.4	11.1	11.9	11.9	8.7	6.5
Middle East and North Africa	9.6	8.0	7.1	5.8	4.8	4.8	5.0	4.2	3.5	2.7
South Asia	61.1	57.4	55.3	53.8	51.7	48.6	45.1	44.3	39.4	36.0
Sub-Saharan Africa	51.5	55.2	54.4	56.5	59.4	58.1	58.0	55.7	52.3	47.5
Total developing world	52.2	47.1	42.3	43.1	40.9	34.8	34.1	30.8	25.1	22.4
Total developing world excluding China	40.5	39.1	38.1	37.2	36.6	34.3	33.6	31.5	27.8	25.2

Source: Chen and Ravallion (2012).

also succeeded in reducing poverty, but at a much slower rate. Excluding China, global poverty declined by less than a third since 1990, from 37.2% to 25.2%.

One notable feature of the global distribution of poverty is that much of it is found outside the poorest countries. For instance, in 2008 India was a lower-middle-income

country,⁴⁴ yet it contained approximately 380 million people below the PPP\$1.25 poverty line,⁴⁵ or 30% of the global total and about the same number as in all of sub-Saharan Africa.⁴⁶ Indeed, Sumner (2012) points out that a majority of people below the World Bank poverty line live in middle-income countries.

In Section 11.6 we noted that within-country inequality has risen over the past two decades, suggesting that the decline in global poverty has been driven by aggregate growth in low- and middle-income countries. It does not follow, however, that continued aggregate growth is the only way to continue to reduce poverty. Redistribution within countries could also play a significant role in poverty reduction in poor countries, just as it does in rich countries.⁴⁷

The average annual income of those living below the PPP\$1.25-a-day poverty line was PPP\$421 in 2005.⁴⁸ Using our estimated global income distribution presented earlier, the individuals in the richest percentile of the world, whose annual incomes averaged PPP\$90,000, were therefore 214 times richer than individuals in the poorest 21% of the world. Put another way, the richest 1% of the world, or 65 million people, had a *total income* a little more than 10 times that of the poorest 21%, or 1390 million people.

11.9. CONCLUSION

Rising top income shares in many countries have pushed inequality up the public agenda, while globalization has given this concern a global reach: people no longer compare their lot only to those within their own country. Moreover, the global financial crisis and recession have made the interconnectedness of people's material well-being around the world all the more obvious. Hence the global distribution of income, global inequality, and global poverty are increasingly in the public view.

⁴⁴ The World Bank defines lower-middle income countries in terms of US\$, not PPP\$. They are defined as those countries with per capita GNI between US\$1036 and US\$4085 in 2012 prices, calculated using the Atlas method. India's per capita GNI in 2008 was US\$1050 in current prices, or US\$1120 in 2012 prices (inflating by the US CPI). In PPP\$, its per capita GDP in 2008 was PPP\$2900 in current prices.

⁴⁵ This is calculated by applying the poverty headcount for 2009 from Povcalnet (the closest year available) to the 2008 population (<http://iresearch.worldbank.org/PovcalNet>). Applying the poverty headcount reported in Table 11.9 to sub-Saharan Africa's population in 2008 implies 389 million poor in that region.

⁴⁶ For this reason Collier's (2007) book *The Bottom Billion*, which is about a set of poor countries whose populations sum to about one billion, is misleadingly titled: those one billion individuals are not the poorest billion people in the world (Segal, 2008).

⁴⁷ In the EU15 countries, for example, 16% of the population were living below their respective national poverty lines in 2003, a figure that would rise to an estimated 25% in the absence of cash benefits that comprise a total of 6.6% of GDP (Guio, 2005). Ravallion (2009) considers the capacity of targeted redistribution within developing countries to eliminate poverty, asking which countries could eliminate poverty by taxing only those with incomes above PPP\$13-a-day. Segal (2011) considers the impact on poverty of a universal unconditional transfer, or "basic income," funded by a country's own natural resource rents.

⁴⁸ This was calculated from a reported poverty gap index of 7.78% in Povcalnet.

A number of studies have estimated the global distribution of income using a variety of data and methods. Recent advances in data collection have provided us with a much more detailed and accurate view of the global distribution of income than was possible even a decade ago. Household surveys now cover the vast majority of the population of the world, whereas intercountry real income comparisons have been greatly improved with the 2005 round of the ICP. This chapter has also highlighted and used the additional information provided by the growing database of top incomes from tax records. Because survey data typically underestimate or underreport the incomes of the very rich, we have estimated global inequality by appending top income data to the available survey data.

This chapter has described the conceptual foundations of the analysis of the global distribution of income, the confusions that can arise by conflating different concepts of that distribution, and the divergent inequality trends that they can display. Its main focus was the global interpersonal distribution of income, or the concept 3 global distribution. Implicitly this analysis assumes a cosmopolitan symmetric social welfare function, according to which the country or location of an individual in the world is irrelevant.

Our calculations show that when the global interpersonal distribution of income is estimated through household survey data *without* top incomes, inequality in this distribution is very high but remains virtually unchanged from 1988 to 2005. The Gini is 0.705 in 1988 and 0.701 in 2005, MLD decreases slightly from 1.063 to 1.060, and Theil T rises marginally from 0.967 to 0.977. However, the income share of the top percentile rises from 11.2% to 14.9%. The equivalent estimates by Milanovic (2012) and Lakner and Milanovic (2013) also find the Gini virtually unchanged at much the same level as us, but the former finds a rise in Theil T, whereas the latter find a fall in both MLD and Theil T.

We argued that the method used in all of these estimates, which take household incomes directly from surveys, was preferable to the method of “scaling” within-country distributions to NA means. Moreover, we argued that if one were to scale in such a way, then HFCE would be preferable to GDP. Bourguignon (2011) scaled within-country distributions to per capita GDP and found a substantial decline in the global Gini coefficient during 1989–2006, which we also find when we use HFCE. However, we find that the divergence in global inequality estimated using HFCE-means and survey-means is due to a reduced coverage of HFCE data relative to survey data in the first year, 1988, and to the sharp divergence between the survey and HFCE means in India in particular.

When we append data on top incomes to the survey distributions, including imputed incomes for countries without such data, we find that inequality is higher, and there is some indication of an increase in global inequality. In this case the Gini, which is less sensitive to inequality at the top end of the distribution, remains virtually unchanged over the period, at a higher level of 0.722 to 0.735. But MLD and Theil T increase over the period—MLD only slightly from 1.136 to 1.156, and Theil T from 1.114 to 1.188. Given the diverse sources of potential error in estimates of the global distribution of income, these changes may not be statistically significant. A larger proportional rise occurs in the income share of the top percentile—from 17.3% to 20.7%. Thus the increase in estimated global inequality

over time with top income data added appears to be driven by the rising income share of the top percentile in the global distribution. We find that in 2005, individuals in the top 1% of the world had an annual average income of PPP\$90,000. These individuals were on average 214 times richer than those in the poorest 21% of the world, who were living below PPP\$1.25-a-day. Put differently, the richest 1% in the world, or 65 million people, had a total income equal to 10 times that of the poorest 21%, or 1.4 billion people.

Inequality in the global distribution can be decomposed into within-country and between-country components using decomposable inequality measures. We find that between-country inequality declined modestly during 1988–2005, whereas within-country inequality increased substantially. Between-country inequality is the larger component of global inequality, comprising 64–81% of overall inequality, depending on the inequality measure (MLD or Theil T) and year. Nevertheless, even if between-country inequality were eliminated, global inequality would remain at about the same level as in a high-inequality country such as China.

Although the extent to which global inequality has risen depends on the measure used, global poverty has declined substantially in recent decades. Given that inequality within countries has tended to rise, this decline has been driven by aggregate growth in low- and middle-income countries. It does not follow, however, that continued aggregate growth is the only way to continue to reduce global poverty. Redistribution of income within countries—by checking or reversing the rise in within-country inequality—could also make a significant contribution to the reduction of global poverty.

ACKNOWLEDGMENTS

We are very grateful to the editors, Tony Atkinson and François Bourguignon, for their detailed comments, and to Branko Milanovic for discussions regarding data. For helpful comments, we would also like to thank Angus Deaton and participants of the “Income Distribution Handbook Conference” at the Paris School of Economics in April 2013.

APPENDIX. ESTIMATES OF GLOBAL INEQUALITY BASED ON THE COMMON SAMPLE OVER TIME

The common sample over time for which we have both survey and PPP data in *all* five benchmark years comprises 67 countries. Table 11.A1 shows the total population of these 67 countries in each year and their share of the world population in that year.

Table 11.A1 Country and population coverage of common sample over time

Year	Number of countries	Population in billions (% of world population)
1988	67	4.16 (82)
1993	67	4.45 (80)
1998	67	4.76 (80)
2002	67	4.98 (80)
2005	67	5.13 (79)

Source: Authors' calculations.

Table 11.A2 Global inequality with and without top incomes, 1988–2005, for common sample over time

Year	Income share of top percentile (%)	Income share of top decile (%)	Gini	Between-country Gini (% of global Gini)	MLD	Between-country MLD (% of total)	Within-country MLD (% of total)	Theil T	Between-country Theil T (% of total)	Within-country Theil T (% of total)
With top incomes										
1988	17.7	58.8	0.730	0.653 (89)	1.159	0.913 (79)	0.246 (21)	1.130	0.796 (70)	0.334 (30)
1993	16.9	57.4	0.726	0.634 (87)	1.155	0.862 (75)	0.294 (25)	1.100	0.747 (68)	0.353 (32)
1998	18.9	59.1	0.722	0.632 (88)	1.093	0.796 (73)	0.297 (27)	1.140	0.754 (66)	0.386 (34)
2002	21.3	60.2	0.729	0.642 (88)	1.115	0.830 (74)	0.286 (26)	1.170	0.781 (67)	0.389 (33)
2005	21.5	58.3	0.721	0.624 (87)	1.141	0.801 (70)	0.340 (30)	1.153	0.722 (63)	0.431 (37)
Without top incomes										
1988	11.4	54.9	0.710	0.646 (91)	1.085	0.886 (82)	0.198 (18)	0.980	0.780 (80)	0.200 (20)
1993	11.6	54.7	0.706	0.630 (89)	1.083	0.843 (78)	0.239 (22)	0.962	0.738 (77)	0.223 (23)
1998	13.6	55.5	0.698	0.624 (89)	1.008	0.772 (77)	0.236 (23)	0.963	0.734 (76)	0.228 (24)
2002	13.3	56.0	0.705	0.633 (90)	1.029	0.798 (78)	0.231 (22)	0.991	0.760 (77)	0.231 (23)
2005	13.2	53.8	0.693	0.611 (88)	1.044	0.767 (73)	0.277 (27)	0.939	0.691 (74)	0.248 (26)

Source: Authors' calculations.

REFERENCES

- Ackland, R., Dowrick, S., Freyens, B., 2004. Measuring Global Poverty: Why PPP Methods Matter. Paper presented at the 2004 Conference of the International Association for Research in Income and Wealth, Cork, Ireland, August. <http://www.iariw.org/papers/2004/dowrick.pdf>.
- Almås, I., 2012. International income inequality: Measuring PPP bias by estimating Engel curves for food. *Am. Econ. Rev.* 102 (1), 1093–1117.
- Alvaredo, F., Londoño Vélez, J., 2013. High Incomes and Personal Taxation in a Developing Economy: Colombia 1993–2010, Commitment to Equity Working Paper 12.
- Anand, S., 1983. Inequality and Poverty in Malaysia: Measurement and Decomposition. Oxford University Press, Oxford.
- Anand, S., Kanbur, S.M.R., 1993. Inequality and development: a critique. *J. Develop. Econ.* 41 (1), 19–43.
- Anand, S., Segal, P., 2008. What do we know about global income inequality? *J. Econ. Lit.* 46 (11), 57–94.
- Anand, S., Sen, A., 2000. The income component of the human development index. *J. Human Develop.* 1 (1), 83–106.
- Anand, S., Segal, P., Stiglitz, J.E., 2010. Introduction. In: Anand, S., Segal, P., Stiglitz, J.E. (Eds.), *Debates on the Measurement of Global Poverty*. Oxford University Press, Oxford.
- Atkinson, A.B., 1970. On the measurement of inequality. *J. Econ. Theory* 2 (3), 244–263.
- Atkinson, A.B., 2007. Measuring top incomes: methodological issues. In: Atkinson, A.B., Piketty, T. (Eds.), *Top Incomes over the Twentieth Century: A Contrast between Continental European and English-Speaking Countries*. Oxford University Press, Oxford.
- Atkinson, A.B., Brandolini, A., 2001. Promise and pitfalls in the use of ‘secondary’ datasets: income inequality in OECD countries. *J. Econ. Lit.* 39 (3), 771–799.
- Atkinson, A.B., Brandolini, A., 2010. On analyzing the world distribution of income. *World Bank Econ. Rev.* 24 (1), 1–37.
- Atkinson, A.B., Piketty, T., Saez, E., 2011. Top incomes in the long run of history. *J. Econ. Lit.* 49 (1), 3–71.
- Banerjee, A., Piketty, T., 2010. Top Indian incomes, 1922–2000. In: Atkinson, A.B., Piketty, T. (Eds.), *Top Incomes: A Global Perspective*. Oxford University Press, Oxford.
- Beddoes, Z.M., 2012. For Richer, For Poorer. *The Economist*. 13 October 2012, <http://www.economist.com/node/21564414>.
- Bernstein, A.R., 2011. Political cosmopolitanism. In: Chatterjee, D. (Ed.), *Encyclopedia of Global Justice*. Springer, London.
- Berry, A., Bourguignon, F., Morrisson, C., 1983. The level of world inequality: how much can one say? *Rev. Income Wealth* 29 (3), 217–241.
- Boltho, A., Toniolo, G., 1999. The assessment: the twentieth century—achievements, failures, lessons. *Oxf. Rev. Econ. Pol.* 15 (4), 1–17.
- Bourguignon, F., 2011. A turning point in global inequality . . . and beyond. In: Krull, W. (Ed.), *Research and Responsibility: Reflection on Our Common Future*. CEP Europäische Verlagsanstalt Leipzig GmbH, Leipzig.
- Bourguignon, F., Morrisson, C., 2002. Inequality among world citizens: 1820–1992. *Am. Econ. Rev.* 92 (4), 727–744.
- Bourguignon, F., Levin, V., Rosenblatt, D., 2004. Declining international inequality and economic divergence: reviewing the evidence through different lenses. *Écon. Intern.* 100, 13–25.
- Burkhauser, R.V., Feng, S., Jenkins, S.P., Larrimore, J., 2009. Recent Trends in Top Income Shares in the USA: Reconciling Estimates from March CPS and IRS Tax Return Data, NBER Working Paper 15320, Cambridge, MA.
- Chen, S., Ravallion, M., 2001. How did the world’s poorest fare in the 1990s? *Rev. Income Wealth* 47, 283–300.
- Chen, S., Ravallion, M., 2008. The Developing World is Poorer than We Thought, but No Less Successful in the Fight against Poverty, World Bank Policy Research Working Paper 4703, Washington, DC.
- Chen, S., Ravallion, M., 2010. China is poorer than we thought, but no less successful in the fight against poverty. In: Anand, S., Segal, P., Stiglitz, J.E. (Eds.), *Debates on the Measurement of Global Poverty*. Oxford University Press, Oxford, pp. 327–340.

- Chen, S., Ravallion, M., 2012. An Update to the World Bank's Estimates of Consumption Poverty in the Developing World. Briefing note prepared by Shaohua Chen and Martin Ravallion, Development Research Group, World Bank, Washington, DC (03-01-12).
- Chotikapanich, D., Valenzuela, R., Prasada Rao, D.S., 1997. Global and regional inequality in the distribution of income: estimation with limited and incomplete data. *Emp. Econ.* 22, 533–546.
- Collier, P., 2007. *The Bottom Billion: Why the Poorest Countries are Failing and What Can Be Done About It*. Oxford University Press, Oxford.
- Deaton, A., 2005. Measuring poverty in a growing world (or measuring growth in a poor world). *Rev. Econ. Stat.* 87 (1), 1–19.
- Deaton, A., 2010. Price indexes, inequality, and the measurement of world poverty. *Am. Econ. Rev.* 100 (1), 5–34.
- Deaton, A., Dupriez, O., 2011. Purchasing power parity exchange rates for the global poor. *Am. Econ. J. Appl. Econ.* 3 (2), 137–166.
- Deaton, A., Heston, A., 2010. Understanding PPPs and PPP-based national accounts. *Am. Econ. J. Macroecon.* 2 (4), 1–35.
- Deaton, A., Kozel, V. (Eds.), 2005. *The Great Indian Poverty Debate*. Macmillan, New Delhi.
- Deininger, K., Squire, L., 1996. A new data set measuring income inequality. *World Bank Econ. Rev.* 10 (3), 565–591.
- Dikhanov, Y., Ward, M., 2002. Evolution of the Global Distribution of Income, 1970–99. Paper prepared for the 53rd Session of the International Statistical Institute held in Seoul, Republic of Korea, August 22–29, 2001 and updated for the 5th Conference on Globalization, Growth and (In)Equality held in Warwick, England, March 15–17, 2002.
- Dowrick, S., Akmal, M., 2005. Contradictory trends in global income inequality: a tale of two biases. *Rev. Income Wealth* 51 (2), 201–229.
- Durlauf, S.N., Johnson, P.A., Temple, J.R.W., 2009. The econometrics of convergence. In: Mills, T.C., Patterson, K. (Eds.), *Palgrave Handbook of Econometrics, Volume 2: Applied Econometrics*. Palgrave Macmillan, New York, NY.
- Firebaugh, G., 1999. Empirics of world income inequality. *Am. J. of Sociol.* 104 (6), 1597–1630.
- Firebaugh, G., 2003. *The New Geography of Global Income Inequality*. Harvard University Press, Cambridge, MA.
- Guio, A.-C., 2005. Income poverty and social exclusion in the EU25. In: *Statistics in Focus, Population and Social Conditions*. Eurostat, 13/2005.
- Harttgen, K., Klasen, S., Vollmer, S., 2013. An African growth miracle? Or: what do asset indices tell us about trends in economic performance? *Rev. Income Wealth* 59 (Special Issue), S37–S61.
- Jerven, M., 2013. Comparability of GDP estimates in sub-Saharan Africa: the effect of revisions in sources and methods since structural adjustment. *Rev. Income Wealth* 59 (Suppl. S1), S16–S36.
- Kolm, S.-C., 1969. The optimal production of social justice. In: Margolis, J., Guitton, H. (Eds.), *Public Economics. An Analysis of Public Production and Consumption and Their Relations to the Private Sectors*. Macmillan, London.
- Korzeniewicz, R., Moran, T., 1997. World-economic trends in the distribution of income, 1965–1992. *Am. J. Sociol.* 102 (4), 1000–1039.
- Lagarde, C., 2013. A New Global Economy for a New Generation. Managing Director, International Monetary Fund, Speech as prepared for delivery, International Monetary Fund, Davos, Switzerland.
- Lakner, C., Milanovic, B., 2013. Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession. World Bank Policy Research Working Paper 6719, December, Washington, DC.
- Melchior, A., Telle, K., Wiig, H., 2000. Globalisation and Inequality: World Income Distribution and Living Standards, 1960–1998. Royal Norwegian Ministry of Foreign Affairs, Studies on Foreign Policy Issues, Report 6B.
- Milanovic, B., 1998. Income, Inequality and Poverty During the Transition from Planned to Market Economy. World Bank Regional and Sectoral Studies, Washington, DC.
- Milanovic, B., 2002. True world income distribution, 1988 and 1993: first calculation based on household surveys alone. *Econ. J.* 112 (476), 51–92.

- Milanovic, B., 2005. *Worlds Apart: Measuring International and Global Inequality*. Princeton University Press, Princeton.
- Milanovic, B., 2012. Global inequality recalculated and updated: the effect of new PPP estimates on global inequality and 2005 estimates. *J. Econ. Inequal.* 10 (1), 1–18.
- Nagel, T., 2005. The problem of global justice. *Philos. Publ. Aff.* 33, 113–147.
- Neary, J.P., 2004. Rationalizing the Penn World Table: true multilateral indices for international comparisons of real income. *Am. Econ. Rev.* 94 (5), 1411–1428.
- Pickett, K., Wilkinson, R., 2010. *The Spirit Level: Why Greater Equality Makes Societies Stronger*. Penguin, London.
- Ravallion, M., 2004. Competing concepts of inequality in the globalization debate. In: Collins, S.M., Graham, C. (Eds.), *Brookings Trade Forum 2004. Globalization, Poverty, and Inequality*. Brookings Institution Press, Washington, DC.
- Ravallion, M., 2009. Do Poorer Countries Have Less Capacity for Redistribution? World Bank Policy Research Working Paper 5046, Washington, DC.
- Reddy, S., Pogge, T., 2010. How *not* to count the poor. In: Anand, S., Segal, P., Stiglitz, J.E. (Eds.), *Debates on the Measurement of Global Poverty*. Oxford University Press, Oxford.
- Rodrik, D., 2013. The Past, Present, and Future of Economic Growth, Global Citizen Foundation. Working Paper 1, http://www.sss.ias.edu/files/pdfs/Rodrik/Research/GCF_Rodrik-working-paper-1_-6-24-13.pdf.
- Sala-i-Martin, X., 2002. The Disturbing ‘Rise’ of Global Income Inequality, NBER Working Paper 8904, Cambridge, MA.
- Sala-i-Martin, X., 2006. The world distribution of income: falling poverty and . . . convergence, period. *Q. J. Econ.* 121 (2), 351–397.
- Schultz, T.P., 1998. Inequality in the distribution of personal income in the world: how it is changing and why. *J. Popul. Econ.* 11 (3), 307–344.
- Seers, D., 1972. What are we trying to measure? *J. Develop. Stud.* 8 (3), 21–36.
- Segal, P., 2008. Review of *The Bottom Billion*, by Paul Collier. *Renewal* 16 (3/4), 152–154.
- Segal, P., 2011. Resource rents, redistribution, and halving global poverty: the resource dividend. *World Dev.* 39 (4), 475–489.
- Sen, A., 1973. *On Economic Inequality*. Clarendon Press, Oxford.
- Sen, A., 2000. Social justice and the distribution of income. In: Atkinson, A.B., Bourguignon, F. (Eds.), *Handbook of Income Distribution, Volume 1*. Elsevier, Amsterdam.
- Sen, A., 2009. *The Idea of Justice*. Harvard University Press, Cambridge, MA.
- Sumner, A., 2012. Where do the poor live? *World Dev.* 40 (5), 865–877.
- Székely, M., Hilgert, M., 1999. What’s Behind the Inequality We Measure: An Investigation Using Latin American Data, Inter-American Development Bank Working Paper #409, Washington, DC.
- Tett, G., 2012. Income Inequality Comes Out of the Igloo at Davos. *FT.com*, 25 January 2012, <http://blogs.ft.com/the-world/2012/01/income-inequality-comes-out-of-the-igloo-at-davos/>.
- World Bank, 1990. *World Development Report 1990: Poverty*. World Bank and Oxford University Press, Washington, DC.
- World Bank, 2008. *2005 International Comparison Program: Tables of Final Results*. World Bank, Washington, DC.
- World Economic Forum, 2012. *Global Risks 2012*, seventh ed. World Economic Forum, Geneva.
- Young, A., 2012. The African Growth Miracle, NBER Working Paper 18490, Cambridge, MA.