

Global Income Distribution

From the Fall of the Berlin Wall to the Great Recession

Christoph Lakner

Branko Milanovic

The World Bank
Development Research Group
Poverty and Inequality Team
December 2013



Abstract

The paper presents a newly compiled and improved database of national household surveys between 1988 and 2008. In 2008, the global Gini index is around 70.5 percent having declined by approximately 2 Gini points over this twenty year period. When it is adjusted for the likely under-reporting of top incomes in surveys by using the gap between national accounts consumption and survey means in combination with a Pareto-type imputation of the upper tail, the estimate is a much higher global Gini of almost 76 percent. With such an adjustment the downward trend in the Gini almost disappears. Tracking the evolution of individual country-deciles shows the underlying elements that drive the

changes in the global distribution: China has graduated from the bottom ranks, modifying the overall shape of the global income distribution in the process and creating an important global “median” class that has transformed a twin-peaked 1988 global distribution into an almost single-peaked one now. The “winners” were country-deciles that in 1988 were around the median of the global income distribution, 90 percent of whom in terms of population are from Asia. The “losers” were the country-deciles that in 1988 were around the 85th percentile of the global income distribution, almost 90 percent of whom in terms of population are from mature economies.

This paper is a product of the Poverty and Inequality Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The authors may be contacted at clakner@worldbank.org and bmilanovic@worldbank.org (or bmilanovic@gc.cuny.edu).

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Global income distribution:
from the fall of the Berlin Wall to the Great Recession¹

Christoph Lakner and Branko Milanovic

JEL codes: D31

Keywords: Income distribution, globalization, top incomes

Sector board: social protection

¹ Both authors are with the World Bank Research Department. The authors would like to thank Statistics Finland, Statistics Portugal and Eurostat for providing tabulated micro data; Maria Ana Lugo and Philippe van Kerm for help with the empirical implementation of the non-anonymous GICs; Shaohua Chen for help with PovcalNet; Tony Atkinson and La-Bhus Jirasavetakul for many helpful discussions. The BHPS data were accessed via the UK Data Service. The paper was in part funded by the World Bank grant under the Knowledge for Change Project (KCP) “Changing inequalities: facts, perceptions and policies” TF012968. Parts of the paper were presented at workshops held at the World Bank and OECD and we thank the participants for useful comments. The opinions expressed in the paper are the authors’, and should not be attributed to the World Bank and its affiliated organizations.

Introduction

This paper provides new evidence on the evolution of global interpersonal income inequality between 1988 and 2008. This inequality concept measures inequality amongst all individuals in the world irrespective of their country of residency thus implicitly assuming a “cosmopolitan” social welfare function and translating the concern for within-country inequality to the global level. Over the period 1988-2008, the face of globalisation changed dramatically with the integration of many developing countries into the world economy. Global interpersonal inequality captures the effects of these shifts on both within- and between-country inequality.

We find that the inequality in the global income distribution, as measured for example by a Gini index, does not change very much over this period. However, this hides substantial re-ranking of country-deciles and changes in the regional composition of different parts of the global distribution. We first present new evidence on the evolution of global interpersonal inequality. We then dig deeper and analyse changes in the composition of the global distribution of income.

Measuring global inequality empirically is substantially more difficult compared with within-country inequality. In the absence of a global household survey, we need to resort to combining national surveys. Our database includes 565 household surveys across five benchmark years and each country-year observation is represented by the average income of ten income decile groups.² National surveys collect information in terms of local currencies, so we need to convert these into a common currency preferably adjusting for differences in the price level (across countries *and* over time).³ We should point out that in constructing our global distribution we mix income and consumption surveys. We refer to them interchangeably, as is customary in this literature, although we are obviously fully aware of the important differences between the two concepts.⁴

² Each decile is weighted by its population, so we measure interpersonal global income inequality where each individual is assigned the income of his or her income decile.

³ Analyses of national income inequality are often done on nominal incomes, thus assuming a common national price level. We follow these national studies by ignoring differences in the price level within countries (except for the case of China, India and Indonesia where we allow for rural-urban price differentials).

⁴ However, we improve on earlier approaches by keeping the type of survey (income or consumption) constant over time for a particular country.

This paper offers four main contributions to the study of global income inequality. First, we compile a new and improved database of national household surveys in response to criticism of earlier data sets (for example by Anand and Segal, 2008). Second, this allows us to present more credible results on global interpersonal income inequality between 1988 and 2008. Third, we create balanced and unbalanced panels of country-deciles for five benchmark years. This allows us to go further than the statements about which countries (and by how much) affected global inequality, by looking into a more disaggregated distribution of country-deciles. We can thus identify those country-deciles that have gained and lost most over this twenty year period. Fourth, we present one of the first comprehensive adjustments for missing top incomes in the study of global inequality.⁵ We offer valuable new empirical insights because the effect on global inequality of non-response at the top within individual countries is unclear on a priori basis (Deaton, 2005).⁶

We estimate the Gini index to be around 70%. The global Gini index has fallen over this twenty year period, with the decline being strongest between 2003 and 2008. However, these observed changes are probably not robust to plausible standard errors. The time-series pattern is robust to alternative measures of inequality, such as the Theil indices. Most of global inequality is accounted for by differences between countries, although this contribution has declined over time, suggesting that countries have become more similar. The within-country component of global inequality, however, has increased continuously over this twenty-year period.

⁵ Pinkovskiy (2013) estimates nonparametric bounds on the global Atkinson index allowing for any country-level income distribution given fractile shares and a Gini index. With a sufficiently high non-response at the top, the direction of change of global *welfare* between 1970 and 2006 becomes ambiguous. He does not present bounds for global *inequality* measures which allow for non-response at the top.

⁶ It might appear intuitive that stretching out the top tail would increase within-country inequality. However, this is not true in general and depends on the particular pattern in which underreporting occurs. As Deaton (2005) shows, if the probability of compliance at the top decreases following a Pareto distribution and the true distribution of incomes is lognormal, the “true” variance will not be different from the one obtained from a truncated distribution. The effect on global inequality depends on the extent of top income misreporting across countries, where in the global income distribution countries and their top income groups are, how far in the national income distribution non-participation begins (e.g. top 10% or top 5%), and how populous the countries and their top recipients are. If top income misreporting is particularly strong in poor countries, global inequality might actually fall once we allow for underreporting at the top.

We present a number of robustness checks. When we scale survey incomes to final household consumption from national accounts, we obtain a lower level of the global Gini index and a stronger decline over time.⁷ We also present a simple robustness check for underreporting of top incomes in household surveys. We treat the discrepancy between national accounts consumption and household surveys, an issue which has received considerable attention in the literature, as a proxy for the extent of “missing” top incomes. We obtain detailed top quantiles of the distribution by allocating this gap to the top decile and fitting a Pareto distribution. The Gini index of this revised distribution is about 5 Gini points higher and does not decline substantially between 1988 and 2008. The difference in levels is primarily due to allocating the national accounts excess consumption to the top decile rather than to the Pareto imputation itself, that is the elongation of the upper tail of the distribution.

The regional composition of the global distribution has changed substantially over this twenty-year period. China has emerged from the bottom ranks of the global distribution, which has had a profound effect not only on the regional composition, but also on the overall shape of the global distribution. Both China’s average income growth and change in income inequality have been exceptionally strong. India has grown more slowly, but its inequality has also grown much less. As a result of the growth in China (and, to a lesser extent, India), Sub-Saharan Africa now makes up most the bottom part of the global distribution.

Not surprisingly, China (particularly the urban part) is among the country-deciles which have grown most between 1988 and 2008. In Latin America, some country-deciles have also performed very well. The new EU member states can be found both amongst the biggest gainers and losers, similarly to Sub-Saharan Africa after 1993.

In constructing our global income distribution we have tried to be as careful as possible given the data constraints and corrected some of the biases in earlier studies (see below). Nevertheless, sources of uncertainty over our estimates remain which are difficult (or impossible in some cases)

⁷ We want to stress that scaling survey incomes to national accounts is not our preferred estimation strategy. We simply replicate the commonly adopted approach in this literature which scales survey incomes to GDP from national accounts. Ours is the first paper to scale to national accounts consumption (rather than GDP), which should be preferred (Anand and Segal, 2008) as we argue in more detail below.

to quantify. We would suggest a conservative approach and conclude that the changes we observe over time are not statistically significant. We return to some of these issues in the conclusion and suggest ways in which we might address them in future work.

This paper is structured in six main parts. First, we provide an overview of the literature on global inequality, the measurement of purchasing power parities (PPPs), the discrepancy between national accounts and household surveys, the underestimation of top incomes in household surveys and outline our approach to all these issues. Section 2 summarises our data construction and methodology, including the welfare concept used and the Pareto imputation to account for top incomes. Section 3 provides summary statistics on our database, particularly the coverage of global and regional GDP and population, and presents the main results regarding inequality in the global distribution of income. We test for the robustness to different measures of inequality and investigate the changing regional composition as well as global growth incidence curves. In Section 4, we provide an upper bound on the global Gini by accounting for the underestimation of top incomes. Section 5 moves from a cross-sectional focus to a panel analysis, taking account of the movement of individual country-deciles in the global distribution. Section 6 presents conclusions.

1 Literature review and our approach

Our paper is related to a number of different strands of the literature. First, we summarise the literature on global inequality, and define what we mean by this term. Second, we explain the various problems associated with deriving PPP exchange rates and why we consider the rates used in this paper to be the most robust. Finally, we address previous work on (a) the underreporting of top incomes in household surveys, and (b) the discrepancy between national accounts measures of income or consumption and their household survey equivalents. We argue that the two issues are closely related and consider them jointly in our analysis.

1.1 Global inequality

Milanovic (2005) distinguishes between three concepts of global inequality. First, unweighted international inequality is the inequality in per capita incomes amongst the countries in the world. Second, population-weighted international inequality, or between-country inequality (Anand and Segal, 2008), measures inequality amongst persons by assigning everybody the per capita income of his place of residence. It thus ignores any within-country inequality. Third, global interpersonal inequality captures the inequality of individual incomes in the world by giving everybody his or her own income.⁸

In this paper we focus on the last concept, global interpersonal income inequality, and whenever we use the term “global inequality” this is concept we refer to. Unweighted international inequality (Concept 1) might be appropriate in studies of income convergence across countries (e.g. Barro and Sala-i-Martin, 1992), but it is not a measure of interpersonal inequality, not least because the weight attached to individuals depends on their country of residence. Bourguignon (2011b) shows that between 1989 and 2006, unweighted international inequality (concept 1) has continued to increase, whereas global inequality (concept 3) has declined, as measured by the P90/P10 ratio. This can be explained by the fact that some populous Asian countries have grown very fast, whereas some smaller (mostly African) countries lagged behind or even declined. Population-weighted international inequality (Concept 2) ignores within-country inequality which seems inappropriate given the widespread concerns on precisely this topic. It might be best

⁸ Anand and Segal (2008) add concept zero, which is the inequality in total (rather than per capita) income among countries.

seen as an intermediate (downward-biased) step towards global interpersonal inequality when survey data for measuring within-country distributions is not available (Milanovic, 2005).

Anand and Segal (2008) distinguish between two reasons why we might be interested in measuring the extent of global inequality. First, out of a concern for “global justice” we might be intrinsically concerned about the distribution of resources amongst the world’s citizens which mirrors the concern for inequality at the country-level (Pogge, 2002, Singer, 2002). This cosmopolitan view of the world assumes a global social welfare function which treats persons the same irrespective of their country of residence (Atkinson and Brandolini, 2010).⁹ Second, changes in global inequality capture some of the effects of globalisation. Over the twenty-year period analysed in this paper, the world economy has become more integrated. We want to stress that any estimate presented here cannot be given a causal interpretation since there exist no counterfactual world economies. In addition, 1988 was certainly not the start of globalisation. However, since then the pattern of global trade and capital flows has changed dramatically, with the integration of China (Haskel et al., 2012), other developing countries (Goldberg and Pavcnik, 2007) and Russia into the world economy. Bourguignon (2012) also considers the effects of globalisation on global inequality, but he focuses on its effects on within-country distributions.¹⁰

Ideally global inequality would be measured from a single representative global household survey, which would be analogous to measuring country-level inequality from national household surveys. In the absence of such a survey, we have to rely on combining national household surveys. Most of the literature on global inequality uses (i) distributional information (e.g. Gini indices) from secondary datasets, such as Deininger and Squire (1996), (ii) assumes that incomes or consumption are everywhere distributed according to a lognormal distribution, and (iii) uses average incomes from national accounts (e.g. Bourguignon and Morrisson (2002) and Sala-i-Martin (2006)). Point (iii) implies a rejection of the often available mean income or consumption

⁹ This world view is not shared by everyone. For example, Bhagwati (2004) calls the concern with global interpersonal inequality a “lunacy”, because the individuals around the world “do not belong to a ‘society’ in which they compare themselves with the others” (p.67). Using a simulated world income tax, Kopczuk et al (2005) find that the current levels of foreign aid are consistent with preferences which value foreigners much less (or with the assumption that most international transfers are wasted).

¹⁰ In addition, our methodologies are quite different, since he uses GDP per capita combined with distributional data from PovcalNet and the OECD.

from household surveys. Points (i)-(iii) are jointly needed to derive income levels at all points of the assumed distribution. According to Anand and Segal (2008), Milanovic (2002, 2005, 2012) is the only author who works directly with the household survey data without scaling to national accounts and we follow this approach in our baseline specification.

Anand and Segal (2008) offer a detailed overview of the literature on global interpersonal inequality to date. All studies agree that the level of inequality is very high, with a Gini index between 63.0% and 68.6% in the 1990s. Because methodologies and data sources differ substantially (e.g. the use of national accounts aggregates, the estimation of within country distributions, the use of different and often inconsistent PPP exchange rates), there exists substantial uncertainty over the direction of change in global inequality. Hence “there is insufficient evidence to reject the null hypothesis of no change in global interpersonal inequality over 1970–2000” (p. 91), to which Pinkovski (2013) agrees using a very different methodology (see above). In our results section we compare our estimates with these previous results.

1.2 PPP exchange rates

Comparing incomes in different countries requires the use of exchange rates. If the law of one price held and there were no non-tradables, we could simply use market exchange rates (Deaton and Heston, 2010). This, however, is clearly not the case and market exchange rates would understate the real standard of living in poor countries, thus overstating global inequality (Anand and Segal, 2008).¹¹ We use PPP exchange rates in order to account for differences in the cost of living across countries. PPP exchange rates convert a given local currency into US\$, the numeraire. Because we are dealing with household income or consumption, we use the PPP exchange rates for private consumption rather than the GDP conversion factors.

The first step in the computation of PPPs involves the collection of price data around the world by the International Comparison Program, which in its most recent round has been coordinated by the World Bank. The latest round of price comparisons refers to year 2005. This round has the

¹¹ Alessandria and Kaboski (2011) explain the failure of the law of one price even for tradables by a model which assumes that poor country consumers have a comparative advantage in search activities.

largest global coverage including 146 countries up from 117 in 1993, the previous round. China participated for the first time ever and India for the first time since 1985.

In addition to the improvements in country coverage, the survey methodology was also improved in the latest round of the ICP, in particular in terms of product specifications. However, the issue of urban bias in price collection (that is, of price data collected disproportionately in urban areas) has received particular attention in the case of China, where the 2005 ICP round led to a substantial upward revision of the previous price level (which had been mostly based on guesswork). The price survey was conducted in 11 metropolitan and periurban areas which were chosen because they were likely to have outlets that sold the products compared in the ICP survey (Chen and Ravallion, 2010a). Chen and Ravallion (2010a) argue that the measured price level is representative of urban prices, but substantially overstates the rural price level. In this paper we follow the approach adopted by Chen and Ravallion (2010b) of treating the official PPP rate as representative of urban China and using a downwardly adjusted PPP rate for rural China.

The second step in the estimation of PPP exchange rates involves the computation of a price index, i.e. a particular weighting scheme which combines the national prices collected in the first step. In the most recent 2005 ICP round, the World Bank has used the index due to Eltetö and Köves (1964) and Szulc (1964) (EKS). The Penn World Tables and earlier estimates by the World Bank used the Geary-Khamis (GK) method (Khamis, 1972). The EKS index is a multilateral price index which combines all the bilateral Fisher price indices.¹² More precisely it is the geometric mean of all the indirect Fisher indices between the base country and the country of interest.¹³ The EKS index satisfies transitivity, so we have one index per country instead of a matrix of indices. But the EKS method violates the “independence of irrelevant country”

¹² The Fisher index is the geometric mean of the Paasche and Laspeyres indices. Its weights thus take into account both the reference and comparison countries. Importantly, it naturally obeys the reversal property, which means that the price level of country 1 based on country 2 is the inverse of the price level of country 2 based on country 1. In the special case of identical homothetic preferences between two countries, the Fisher index is a second-order approximation to a “true” cost-of-living index (Deaton and Heston, 2010).

¹³ Suppose we are interested in estimating Ghana’s price level (relative to the US, the base country). There exists a Fisher index which goes from the US to Germany and a separate Fisher index which goes from Germany to Ghana. The EKS index is a geometric mean of all these indirect Fisher indices.

property, because the index between any two countries is affected by prices and weights in third countries (precisely because it combines the indirect Fisher indices).

The GK index compares domestic prices with world prices. The problem is that in the computation of these world prices, the weight attached to a particular country depends on its physical volume of consumption. Thus in practice, the rich countries dominate these composite world prices. Since goods that are relatively expensive in rich countries (say, services) tend to be consumed in relatively large quantities in poor countries, precisely because they are cheaper, the use of a GK index tends to overestimate the value of consumption in poor countries. This is simply a manifestation of the well-known Gerschenkron (1947) effect (or substitution bias) which says that a country's consumption is overvalued when evaluated at the prices of another country, and the further the two price vectors, the greater the overvaluation. The EKS index does not suffer from this bias because it averages the consumption weights from both countries, making "a compromise that is arguably the best that can be done in the circumstances" (Deaton and Heston, 2010, p.11).¹⁴ As a result of the Gerschenkron effect, the GK index understates global inequality (and global poverty, see Ackland et al. (2013)). Deaton and Heston (2010) compute population-weighted between-country inequality of GDP per capita using different indices. They obtain a Gini index of 53.3% for EKS and a value of 52.7% for GK.

In summary, we use the EKS index as suggested by Anand and Segal (2008), Deaton and Heston (2010), Ravallion (2010), and Ackland et al. (2013). For other purposes, the GK index which satisfies additivity might be more relevant.¹⁵ We use a single PPP exchange rate per country (differentiating only between rural/urban China, India and Indonesia), thus ignoring any consumption and price differences along the income distribution.¹⁶

¹⁴ For example, consider alcohol in Bangladesh, where it has a small share, but a relatively high price (see Deaton and Heston, 2010). The Fisher index strikes a compromise between two extreme positions: Bangladeshi budget shares would understate Bangladeshi prices, whereas OECD shares would overstates them.

¹⁵ For example, this might be important when studying the composition of GDP because components converted at GK conversion factors are additive which is not the case with the EKS index. This is a reason why the Penn World Tables continue to use the GK approach.

¹⁶ Reddy and Pogge (2010) argue that consumption PPPs should not be used in the measurement of global poverty because the consumption baskets of the poor are systematically different from the rest of the population. For example, the poor might face different prices because of where they buy (Frankel and Gould, 2001) or how much they buy (Rao, 2000). However, Deaton and Dupriez (2009) find that re-weighting commodity baskets to account for

The final issues in the use of PPPs are how to extend them over time. We compare prices across countries only once using the most recent ICP round, relying on domestic consumer price inflation for the within-country comparisons. In other words, our approach only requires one ICP round: domestic local currency units in any year are converted into domestic 2005 prices using a domestic CPI deflator, and to these (constant price) local currency units we then apply the 2005 PPP exchange rate obtained from direct price comparisons under ICP.

Conceptually, our approach is simple because it keeps the comparisons over space and over time separate. All our within-country comparisons are independent of international prices and only depend on domestic prices, which is attractive not least because domestic prices are appropriate for assessing the trade-offs at the country-level (Nuxoll, 1994).

1.3 National accounts and household surveys

Typically, per capita household consumption in national accounts exceeds the average consumption or income recorded in the survey (Deaton 2005).¹⁷ Moreover, as Deaton shows, the discrepancy appears to have increased over time not only in India (which is somewhat of a *cause célèbre* in that respect) but also in rich countries like the United States and Great Britain. Studies of global inequality differ in their view of how to account for this discrepancy. In our main specification, we follow the approach suggested by Anand and Segal (2008), and simply use the average income recorded in the survey.

As mentioned before, other papers (except Milanovic, 2002, 2012) have anchored the income level to the national accounts (usually, GDPs per capita), combined it with distributional information from household surveys and typically assumed lognormality. Anand and Segal (2008) argue that GDP per capita “is not a suitable measure of household income” (p. 67) and should not be used to anchor household survey means, since it includes items such as

different consumption baskets of the poor and non-poor has very little effect on the estimated PPP exchange rates, which is also reassuring for our estimation strategy.

¹⁷ In some - mostly African - countries, per capita national accounts consumption is actually lower than that found in the household survey. Deaton (2005) argues that this might be explained by under-estimation in the national accounts rather than by problems in the household surveys. For an early statement of the issue, see Milanovic (2002, pp. 65-6).

depreciation, retained corporate earnings or taxes which are not distributed back to households, all of which are only remotely related to household income. Furthermore, there exists “a basic incongruity in assuming that the relative within-country distributions are measured acceptably well by surveys but their means are not” (Anand and Segal, 2008, p. 70). In addition, the replacement of the survey mean by a typically larger GDP per capita implies an equi-proportional adjustment of all incomes. This type of adjustment, which we call “proportional”, is very unlikely to be correct because it implies the same income underestimation, in relative terms, across the entire distribution.

Compared with GDP, final household consumption expenditure from national accounts is closer to household income (or consumption) recorded in surveys (Anand and Segal, 2008).¹⁸ However, it must be noted, that the data and methods used to estimate national accounts consumption are not necessarily more reliable than household surveys (Anand and Segal, 2008, Deaton, 2005).¹⁹ There are also definitional differences between household surveys and national accounts consumption, such as the inclusion of the imputed value of owner-occupied housing (although conceptually it should be included in both but in practice it often is not included in household surveys), imputed financial services, and consumption by non-profit institutions (Deaton 2005).

1.4 Top incomes

Recent work on top share inequality using tax records argues that top incomes are understated in standard household surveys. This literature studies inequality at the very top of the distribution, typically expressed in top shares, e.g. the share of total income received by the top 1% (Atkinson and Piketty (2007, 2010)). Tax data might provide more accurate information on top income recipients for a number of reasons: First, it might be harder to enter the gated communities of the rich than to conduct surveys in poor areas, so survey non-response would increase with income (Groves and Couper, 1998). Second, the top 1% are rare by definition, so a household survey with a standard sample size of a few thousand would offer top share estimates with a low

¹⁸ It might be argued that *income* recorded in household surveys should be approximated with GDP (rather than national accounts consumption). However, Anand and Segal (2008) argue that even in this case, national accounts consumption is to be preferred because of the unrelated components included in GDP mentioned above.

¹⁹ The measurement issues in national accounts include the measurement of illegal transactions or the measurement of intermediate inputs. Furthermore, because consumption is computed as a residual, measurement errors are compounded (Anand and Segal, 2008).

precision, or might miss these people altogether. On the other hand, the tax data intentionally oversample the rich. Third, aspects of survey design such as top-coding or the elimination of “outliers”, manipulate top incomes. On the other hand, tax data are not without problems, e.g. due to tax evasion and income minimisation which may be particularly strong in developing countries.

There is some evidence to support the argument that top incomes are missing in household surveys. Alvaredo (2010) finds that a household survey in Argentina records no observations with incomes exceeding \$1 million whereas the Argentinian tax data contain close to 700 observations in that range. In a comparison of household surveys from 16 Latin American countries, Székely and Hilgert (1999) find that the 10 richest households in the survey receive incomes similar to a managerial wage. It would appear plausible that the top capital owners in these countries receive substantially greater incomes than a manager. Some studies compare top shares estimated from household surveys and tax data, and in some cases obtain very similar results, although this typically depends on the availability of exceptional surveys which have sufficient sample sizes and are not subject to top coding (Burkhauser et al (2012) using internal United States CPS data, Leigh and van der Eng (2009) for Indonesia, and Morival (2011) for South Africa).

Given that tax data appear to be more accurate at measuring top incomes and household surveys offer more precise information about the rest of the distribution, a natural next step would be to combine the two sources of information to obtain a complete distribution of income. However, the still sparse availability of tax data across countries, limits the usefulness of such an exercise for the purpose of analysing the global distribution. In addition, the population and the welfare measure are fundamentally different between the two data sources which makes such an exercise difficult.²⁰

²⁰ In the tax data, the unit of analysis is a tax unit, which depending on the jurisdiction could be a married couple or an individual. A household would typically be bigger than a tax unit. The tax data literature uses taxable (and usually before-tax) income, whereas the household survey collects disposable income. Taxable income excludes some real income, such as tax-exempt interest on government bonds, and deductions and exemptions, although most empirical work using tax data adds these back. On the other hand, household surveys typically measure capital incomes and gains poorly compared with tax data, which cover at least taxable capital incomes and gains. Because the tax data typically do not contain sufficient information in order to construct units and incomes which are similar to those in a standard household survey, the only possibility is to construct tax units and taxable income in the household survey, as in Alvaredo (2010). In the final step of such an exercise, one needs to assume which parts of the true distribution

1.5 Addressing jointly top income underreporting and the national accounts discrepancy

The underreporting of top incomes in household surveys and their discrepancy with national accounts are closely connected issues. It is reasonable to expect, and there is some empirical evidence to corroborate it, that the discrepancy between surveys and national accounts is not distribution neutral and is largely due to non-participation of the rich in household surveys (Mistaenen and Ravallion 2003; Korinek et al 2006).²¹ Deaton (2005) points out that because national accounts consumption tracks money rather than people, national accounts data are more likely to capture large transactions. Using Indian tax record data, Banerjee and Piketty (2010) find that a significant part of the discrepancy between consumption growth in national accounts and household surveys can be accounted for by underreporting of the rich. Finally, it could be argued that household surveys offer a good approximation to the bottom 90% of the distribution (thus, however, ignoring any underreporting of incomes among the very poor).²²

In the second part of the analysis we allocate the gap between household final consumption in national accounts and household surveys to the top 10% of the distribution and obtain more disaggregated top quantiles by fitting a Pareto distribution to the upper tail. Our approach builds on Atkinson (2007) who uses a Pareto imputation in combination with the Bourguignon and Morrisson (2002) data. Atkinson uses GDP per capita, thus spreading the discrepancy between national accounts and household surveys *evenly* across the distribution, but for the very top “elongates” the distribution by using a Pareto interpolation.²³ We call Atkinson’s approach the “proportional adjustment with Pareto tail”. By contrast, our methodology proposes to allocate the “excess” consumption recorded in the national accounts only to the top decile and to use a Pareto

are represented by the tax data and which are covered by the household survey. CBO (2012) matches US tax records with household survey records using income. It adds the non-taxable income (e.g. transfers or in-kind income) from the survey. Armour et al. (2013) also match records by income but they add capital incomes to the household survey because these incomes tend to be poorly measured in this type of data.

²¹ This can also explain why the discrepancy is increasing in countries such as China or India as they become richer (Anand et al 2010).

²² The inclusion of the poor may be insufficient because of the very definition used by surveys, such that they exclude the homeless and institutionalized populations (see Carr-Hill, 2013) or because of sampling issues (excluding remote and probably poorer regions). Thus the very bottom of the distribution may be truncated. But, in addition, incomes of the poor included in the survey may be mismeasured due to extensive home production or benefits received in kind which may not be always included.

²³ The Pareto imputation does not add “new” observations, but rather “stretches out” the top decile. This implies that the imputation by itself does not change the mean.

interpolation, thus both increasing the mean and changing inequality. We call this adjustment “top heavy adjustment with Pareto tail.”

We justify the adjustment of the survey mean to national accounts consumption on the grounds of missing top incomes. This is however open to criticism. It might be argued that some elements of national accounts consumption should (1) be excluded altogether, or (2) spread more widely along the distribution than the top 10%. For example, some of the discrepancy between national accounts consumption and household surveys is related to differences in definition, such as the inclusion of expenditures by non-profit institutions serving households (NPISH) or imputed consumption of collectively provided goods.²⁴ We ought to subtract these components from the national accounts consumption but sufficiently detailed data are not available separately for a large number of countries. Other sources of the discrepancy, such as imputed rents of owner-occupiers in the national accounts, could be included (if they are not estimated by household surveys), but should be spread further along the distribution than just the top 10%.²⁵ For these reasons our estimates should be seen as an approximate first step, in the absence of a more careful analysis using unit-record data.

²⁴ According to Deaton (2005), NPISH consumption accounted for 3.9% of total consumption in the UK in 2001, up from 2.1% in 1970. In addition, he argues that this share might be even higher in poor countries, although there exist no data.

²⁵ Deaton (2001) argues that in India approximately half of the gap between national accounts consumption and household surveys is due to imputed rents.

2 Data construction and methodology

2.1 Data sources

The data used in this paper consist of country-year average decile income/consumption covering the period 1988 to 2008. This means average per capita income for a given decile in country i and year t . The data come from a number of sources. PovcalNet is the starting point of our database, contributing more than two-thirds of the surveys.²⁶ PovcalNet is the compilation of a large number of household surveys stored by the World Bank research department. It has been mostly used to compute estimates of world poverty, as in Chen and Ravallion (2010b), and thus lacks data on rich countries. From PovcalNet we obtain average per capita incomes, already converted in 2005 \$PPPs, and decile shares, which we combine to compute decile average incomes.²⁷

Next, we merge the updated World Income Distribution (WYD) data (Milanovic 2012). PovcalNet and WYD provide almost 98 percent of all data. We convert these data into country-year deciles in order to obtain a consistent database.²⁸ Where possible we fill remaining gaps with data from the Luxembourg Income Study (LIS), the British Household Panel Survey (BHPS), the European Union Survey of Income and Living Conditions (SILC) and data from country statistical offices.²⁹ Overall we end up with 565 surveys across the five benchmark years 1988, 1993, 1998, 2003 and 2008 (Table 1).

²⁶ PovcalNet is the on-line tool for poverty measurement developed by the Development Research Group of the World Bank, <http://iresearch.worldbank.org/PovcalNet>. Data downloaded on 29 July 2012, which refers to the last data update on 28 February 2012.

²⁷ PovcalNet uses grouped data derived from household surveys to derive these decile shares. They are estimated from Lorenz curves fitted to the population-weighted (accounting for household size and sampling weights) distribution of per capita household income or consumption. Both Generalised Quadratic and Beta Lorenz curves are estimated and the functional form with the better fit is chosen.

²⁸ The vast majority of country-year observations are already in deciles or in equally spaced quantiles (e.g., ventiles), so they can be easily converted. In total 13 country-years were imputed by fitting a log-normal Lorenz curve using the “ungroup” command included in the DASP Package (Abdelkrim and Duclos, 2007). This procedure implements the Shorrocks and Wan (2008) adjustment thus ensuring that the fitted Lorenz curve matches the original points. We choose a log-normal functional form and fitted the Lorenz curve on 2000 observations, as suggested by Shorrocks and Wan (2008). Minoiu and Reddy (2012) show that for estimating the global income distribution it is better to fit a parametric Lorenz curve than to estimate the kernel density.

²⁹ The sources for the final database (country-years in parentheses) are: PovcalNet (379), WYD (173), LIS (8), SILC (2), and one survey each from BHPS (Bardasi et al, 2012), Statistics Finland, and Statistics Portugal.

Each country's distribution is represented by the average incomes of the ten deciles. This is not dissimilar from other studies in the literature such as Bourguignon and Morrisson (2002) who use 11 quantiles. This ignores any within-decile inequality, thus, as argued by Anand and Segal (2008) understating within-country inequality and perhaps global inequality. Our choice of decile groups was dictated by PovcalNet, where more detailed information is unavailable.³⁰ Also, because of consistency and ease of discussion, we decided to use deciles also in those surveys where more detailed information was available.

2.2 Survey selection

The surveys included in the database need to meet two conditions: First, they need to be within two years of a benchmark year. Second, they need to be at least three and no more than seven years from the previous and next survey. The rationale of the second condition is not to allow surveys that are either too close or too far apart from the interval of five years since a lot of the analysis is based on the assumption that five-year intervals hold throughout the sample. Table 1 shows the years between the survey year and the benchmark year.³¹ The choice of benchmark years is essentially arbitrary and we followed Milanovic (2012) in choosing 1988, 1993 and 1998. Global household survey coverage is very poor prior to 1988. 2003 and 2008 were chosen in order to obtain equally spaced benchmark years. Compared with Milanovic (2012), we managed to obtain a closer fit to the benchmark year in all years with roughly $\frac{3}{4}$ of surveys conducted within one year of the benchmark year.

We use a mix of income and consumption surveys, as is customary in this literature. Although there are obviously important differences between income and consumption, we refer to them, as already mentioned, interchangeably.³² We do not adjust for differences between income and consumption surveys because any such adjustment, applied to deciles, would be arbitrary.³³

³⁰ Using PovcalNet, Segal (2011) obtains more detailed fractiles (limited to represent at most 5 million people) by accessing the detailed estimation code of the Lorenz curves which PovcalNet fits on the grouped data. This is cumbersome and was not feasible given the large number of surveys we deal with, but may be addressed in future work.

³¹ For example, for the benchmark year 2008, 7.4% of surveys were conducted in the year 2006. In the case of PovcalNet, the survey year is sometimes not an integer, which would happen if a survey was conducted over more than one year. In all cases we took the start year of the survey, i.e. we rounded down the non-integer years.

³² Inequality tends to be lower in terms of consumption than in terms of income (Deininger and Squire, 1996), and (at least in the USA) has increased by less in consumption-terms (Fisher et al., 2013). In the full PovcalNet data (not the sample used in this paper), the average Gini index of consumption surveys (37.98%) is approximately 10 Gini points

Table 1: Sample summary statistics						
	Benchmark year					Total
	1988	1993	1998	2003	2008	
Number of surveys	75	115	121	133	121	565
<u>Years between survey year and benchmark year (% , by benchmark year)</u>						
-2	12.0	9.6	9.1	7.5	7.4	8.9
-1	26.7	18.3	14.9	18.8	11.6	17.4
0	29.3	34.8	41.3	30.1	65.3	40.9
1	18.7	20.9	18.2	21.1	11.6	18.1
2	13.3	16.5	16.5	22.6	4.1	14.9
Within +/- 1 of benchmark	74.7	73.9	74.4	69.9	88.4	
<u>Income vs. Consumption surveys (% , by benchmark year)</u>						
Consumption	33.3	46.1	48.8	57.1	55.4	49.6
Income	66.7	53.9	51.2	42.9	44.6	50.4
<u>GDP (% of regional GDP represented in the database)</u>						
World	90.6	97.0	96.5	95.9	93.0	94.6
Mature economies	95.7	99.9	99.8	98.4	96.9	98.1
China	100.0	100.0	100.0	100.0	100.0	100.0
India	100.0	100.0	100.0	100.0	100.0	100.0
Other Asia	90.5	96.1	98.0	97.3	97.1	95.8
M. East & N. Africa	52.1	55.4	45.4	48.5	22.3	44.7
Sub-Saharan Africa	21.9	82.9	78.3	81.8	77.4	68.5
L. America & Caribbean	94.8	98.3	99.0	98.9	98.4	97.9
Russia, C. Asia, SE Europe	50.7	94.2	94.3	100.0	90.8	86.0
<u>Population (% of regional population represented in the database)</u>						
World	81.1	92.3	91.9	93.6	90.6	89.9
Mature economies	95.0	99.9	99.6	96.7	97.0	97.6
China	100.0	100.0	100.0	100.0	100.0	100.0
India	100.0	100.0	100.0	100.0	100.0	100.0
Other Asia	74.6	85.4	88.7	88.7	88.7	85.2
M. East & N. Africa	60.4	69.5	63.6	68.4	47.8	61.9
Sub-Saharan Africa	28.5	72.9	68.0	80.2	74.1	64.7
L. America & Caribbean	88.2	92.9	94.9	96.4	94.4	93.4
Russia, C. Asia, SE Europe	28.4	87.4	87.4	99.4	84.4	77.4
<i>Notes: Observations are weighted by population size in the computation of global coverage, otherwise unweighted. The last column is the (unweighted) average over the 5 benchmark year values</i>						

lower than the average Gini over the income surveys (48.44%). This is more than the Gini adjustment proposed by Li et al (1998) of 6.6 Gini points.

³³ For example, the income data could be scaled down by savings in the national accounts (i.e. the gap between national consumption and national income) (Deaton, 2001). We did not adopt such an approach because of (1) the issues with national accounts data discussed above and (2) evidence that the savings rate is not invariant with income (Dynan et al., 2004). Chen and Ravallion incorporate such an adjustment in early estimates of global poverty but abandon it in later work (Chen and Ravallion, 2004).

One of the innovations of our data base is that we restrict the income concept to be the same over time for a given country. This avoids any spurious changes arising from a change in the welfare concept being used.³⁴ For each country, income or consumption was chosen so as to maximise the number of benchmark years covered (subject to the two conditions in the previous paragraph).³⁵ As Table 1 shows, in the overall sample the number of consumption and income surveys is almost equal. In earlier years, the majority of surveys collected information on income, whereas in recent years the reverse is true. This can be explained by the improved survey coverage of poor countries where consumption surveys are more common (with the exception of Latin America).³⁶

2.3 Welfare concept

We are interested in analysing the global distribution of (annual) per capita income (in 2005 \$PPP). Per capita incomes ignore any economies of scale in household consumption and within-household inequality. Per capita incomes have the advantage that they are simple to compute and have natural counterparts in the national accounts (which do not compute equivalised incomes). The effect of using a different equivalence scale on world inequality is not clear a priori.³⁷

In our database, each country-year distribution of per capita income is represented by the average incomes of the ten deciles.³⁸ In the analysis, each decile is weighted by its population (i.e. 10% of

³⁴ Income and consumption inequality might move in different directions over the same period of time (e.g. Krueger and Perri, 2006).

³⁵ For Bulgaria, Botswana and Croatia the number of benchmark years with consumption and income information were the same. In all cases we chose the type of survey in order to maximise the number of surveys drawn from PovcalNet. For Nicaragua, we had both types of surveys in all years from PovcalNet. We chose income surveys since this is the prevalent type of survey in Latin America.

³⁶ In China, India and Middle East/North Africa, our database only uses consumption surveys (with the exception of All-China where income surveys are used). In Sub-Saharan Africa 98% of surveys are consumption surveys. In other Asia, 91% are consumption surveys. On the other hand, in the mature economies and Latin America, 97% and 96% respectively are income-surveys.

³⁷ In their study of the LIS data, Atkinson et al (1995) find that the inequality of per capita household income is greater than the inequality of household income adjusted by a square root equivalence scale. The precise effect on cross-country comparisons of inequality depends on the joint distribution of family size and income.

³⁸ For Switzerland in 2008, the bottom decile had average income of zero. We have recoded this to missing, because otherwise the bottom 10% in Switzerland would be the poorest people in the world. Hence we have one missing observation in 2008.

the national population from the World Development Indicators (WDI)). Whenever we are interested in the performance of our estimation or database, e.g. the split between consumption and income surveys or the size of the Pareto constants, observations are unweighted, as pointed out in the tables.

We use consumption PPPs to account for price differences across countries. Incomes obtained from PovcalNet are already converted to 2005 PPP dollars. For the additional surveys we replicate the approach in PovcalNet: as explained before, and after accounting for currency conversions³⁹, we convert the average incomes into local currency units in 2005 prices using domestic consumer price indices (CPIs).⁴⁰ Then, we apply the 2005 PPP consumption exchange rates to convert into international dollars.⁴¹

It is important to note that PPP exchange rates only exist at the country-level, so we ignore any price differences which exist within countries. As a result we probably over-state within-country inequality. As mentioned before, we treat the rural and urban areas of the three most populous developing countries China, India and Indonesia as separate “countries”. Due to a lack of disaggregated data, we assume a common CPI for rural and urban areas, but allow for different PPP exchange rates.⁴²

The vast majority of surveys cover the entire country except for several, mostly Latin American, countries which survey only urban areas.⁴³ We treat these surveys as representative of the entire country.⁴⁴

³⁹ Currency conversions include changes in the currency being used, such as the formation of the Eurozone, and currency redenominations as often observed in high-inflation environments.

⁴⁰ As a first step, we use domestic CPI figures from the WDI. Where these are unavailable, we resort to the IMF’s World Economic Outlook or the country’s statistical offices directly.

⁴¹ The WDI does not provide PPPs for Kosovo, so we used the implied PPP GDP conversion factors reported in the IMF’s World Economic Outlook.

⁴² Data for China and Indonesia are provided by PovcalNet so they already incorporate adjustments for differential costs of living in rural and urban areas. For India, where we added one survey from the WYD, we use urban and rural PPP exchange rates of Rs 17.24 and Rs 11.40 respectively which are given in Ravallion (2008).

⁴³ These countries include Argentina, Colombia, Ecuador, Honduras, Micronesia and Uruguay, and are all taken from PovcalNet.

2.4 Definition of regions

We have grouped countries into eight regions. The first group consists of “mature economies”, which are the EU-27 countries (members in 2008) plus the high-income countries in the world.⁴⁵ We treat India and China as regions in their own right. The remaining groups are defined as residuals according to the geographic regions used in the WDI.

2.5 Pareto imputation and scaling to national accounts consumption

We allocate the excess of national accounts consumption over household surveys in two steps. First, we adjust the country-mean to equal the maximum of the survey mean and national accounts consumption.⁴⁶ Second, we re-compute the decile shares for all deciles except the top using the original average decile incomes and the adjusted mean (the share in total income of those deciles therefore decreases). Third, we compute the new top decile share as the difference between 100% and the sum of the revised shares of the bottom 9 deciles. We use the revised top 10% and top 20% shares in the Pareto imputation.

We obtain household final consumption expenditure⁴⁷ (in 2005 \$PPP) from the World Development Indicators for the survey years.⁴⁸ It is important to note that the sample changes in this part of the analysis because of two reasons. First, due to the unavailability of macro data, we lose some country-year observations. Second, because of a lack of disaggregated macro data, we

⁴⁴ This is likely to understate within-country inequality in these countries, because we expect the rural areas to be poorer compared with the urban areas. The share of the rural population in several of them (Argentina, Uruguay) is minimal though.

⁴⁵ To be precise, the mature economies include EU-27, Australia, Bermuda, Canada, Hong Kong, Iceland, Israel, Japan, Korea, New Zealand, Norway, Singapore, Switzerland, Taiwan and USA.

⁴⁶ In the majority of cases, national accounts consumption exceeds the income recorded in the survey, so the revised mean equals the national accounts consumption. When the survey mean is greater than per capita national accounts consumption, we keep the survey mean.

⁴⁷ It is defined as “the market value of all goods and services [...] purchased by households” (<http://data.worldbank.org/indicator/NE.CON.PRVT.PP.KD>). It includes durable products, imputed rent for owner-occupiers, payments to obtain permits and licenses, and expenditures of non-profit institutions serving households.

⁴⁸ We filled any gaps in the WDI with data from the IMF’s International Financial Statistics (IFS) or the country’s statistical offices. We complement the series in constant 2005 \$PPP with information from the WDI and the IFS in current and constant local currencies as well as current USD. We convert the current USD using market exchange rates. In the rest of the conversion we follow the same approach as with the micro data using consumption PPPs and CPI.

use the whole-country distributions in the case of China, India and Indonesia (where we previously used separate rural/urban distributions). In the case of China, we now use income surveys where we previously used consumption surveys.⁴⁹

We use a Pareto imputation to split the top decile into smaller quantiles. We choose to split the top 10% into P90-P95, P95-P99 and P99-P100. The resulting data set thus consists of 12 (uneven) fractile groups per country (which are weighted by population in the analysis). The implicit assumption is that the top decile of our database follows a continuous Pareto distribution. Let H_i be the cumulative population share of individuals with incomes greater or equal to y_i , e.g. these might be the top 10%. Let S_i be the share of total income received by this group. Atkinson (2007) shows that for the Pareto distribution, the relative share of two top groups is given by

$$\log\left(\frac{S_i}{S_j}\right) = \left(\frac{a-1}{a}\right) \log\left(\frac{H_i}{H_j}\right) \quad (1)$$

We can re-arrange this to compute the Pareto coefficient a

$$a = \frac{1}{1 - \frac{\log(S_i/S_j)}{\log(H_i/H_j)}} \quad (2)$$

We use the top 10% and top 20% shares to compute the Pareto coefficient a for every country-year observation.⁵⁰ Next we compute the top 1% and top 5% shares by using this estimate of a and solving equation (1) for S_i . Then we can easily construct the new quantile groups. P99-P100 is simply the top 1%. P95-P99 is the top 5% share minus the top 1% share. Finally, P90-P95 is the top 10% share minus the top 5% share. For each country-year, we thus have 12 income fractiles.

The validity of our results obviously rests on the parametric assumptions. Our choice of functional form is relatively standard in the literature, where it is argued that top tails are approximately Pareto. Furthermore, the estimation is relatively flexible, since we estimate a different Pareto constant for each country-year observation.

⁴⁹ For the latest benchmark year, PovcalNet does not have a consumption survey, so we used an income survey from the WYD.

⁵⁰ Atkinson (2007) uses the top 5% and top 10% shares. Hlasny and Verme (2013) use the top 10% and top 1%. The World Top Incomes Database (Alvaredo et al., 2013) uses top 0.1% and top 1% in most cases.

2.6 The panel dimension of our data

We are also interested in changes of a given country-decile over time. Hence, the panel dimension of our data is crucial. Table 2 counts the number of countries by the number of benchmark years for which a country appears in the data. For example, for 58 countries (out of a total of 162), mostly mature economies and Latin American countries, we have the complete panel. When we consider changes between two benchmark years, we do not require observations in the years in between, i.e. there could be gaps in the panel. For 63 countries, we can thus consider changes between 1988 and 2008. As a robustness check, we also consider the period 1993 to 2008, for which we have 90 countries and particularly improve the regional coverage of Sub-Saharan Africa and Russia/Central Asia/South-East Europe.

Regions	Total	No. of benchmark years					No. of countries with data in...	
		1	2	3	4	5	1988 & 2008	1993 & 2008
World	162	22	24	27	31	58	63	90
Mature economies	39	0	2	2	10	25	29	34
China	2	0	0	0	0	2	2	2
India	2	0	0	0	0	2	2	2
Other Asia	19	2	5	1	3	8	8	11
M. East & N. Africa	11	4	1	1	3	2	2	3
Sub-Saharan Africa	43	11	9	14	5	4	4	16
L. America & Caribbean	26	4	3	0	5	14	15	17
Russia, C. Asia, SE Europe	20	1	4	9	5	1	1	5

Notes: The last two columns allow for gaps in the panel.

3 The cross-sectional distribution over time

3.1 Summary statistics

Since we are interested in analysing the world distribution of income, a first question to ask is how much of the world is represented by the surveys included in our database (Table 1). Because high-income countries are more likely to have a survey which can be included in our data, our coverage is higher when measured in terms of GDP than in terms of population. Our data represent 95% of world GDP on average and more than 90% in all benchmark years. On average (and in all years since 1993), our data also cover 90% of the world's population.

There are, however, substantial differences across regions. The coverage of Sub-Saharan Africa and Russia/Central Asia/SE Europe has improved markedly, in particular after 1988. Our coverage of the Middle East and Northern Africa region, appears to have declined, particularly in the most recent benchmark year and more so in terms of GDP than population.⁵¹ In the latter part of the analysis we focus on the period from 1993 to 2008, because 1988 has such a poor coverage of Sub-Saharan Africa and Russia/Central Asia/SE Europe.

3.2 The Gini index of the global distribution of income

Table 3 presents our main results on the inequality in the global distribution of income calculated across the unbalanced panel of country-deciles. Compared with within-country distributions, we find a very high level of inequality as measured by the Gini index: between 70.5% and 72.2%.⁵² The global Gini index has virtually remained unchanged. Changes between benchmark years have been around 0.5%, with the exception of the period 2003 to 2008, when the Gini decreased by 1.89% or 1.35 Gini points. The Lorenz curves for 1988 and 2008 (not shown here) intersect.

⁵¹ This appears to be driven by the dropping out of Iran and Tunisia in 2008, which together represent 26% of the region's GDP and 23% of the region's population in 2008. Coverage in this region remains low because we miss big countries such as Saudi Arabia or the United Arab Emirates which account for 17% and 10% of regional GDP respectively.

⁵² For example, the Gini indices at the country-level reported in PovcalNet (the full sample, not the sample used in this paper) range from 19.4% to 74.3%, with an average of 42.2%. Only Jamaica (70.81%) and Namibia (74.33%) have Gini coefficients exceeding 70%. This dataset excludes rich countries, which tend to have lower inequality, so the average Gini is probably upward-biased.

	Benchmark year					1988-2008 change (%)	1993-2008 change (%)
	1988	1993	1998	2003	2008		
Global inequality							
Gini index (%)	72.2	71.9	71.5	71.9	70.5	-2.3	-2.0
GE(0) (Theil-L) (%)	114.2	110.7	107.1	107.6	102.7	-10.1	-7.2
GE(1) (Theil-T) (%)	102.2	102.4	102.8	104.9	100.3	-1.9	-2.1
GE(2) (%)	173.7	179.2	193.0	204.3	201.4	15.9	12.4
Atkinson index A(2) (%)	83.5	82.8	81.8	82.0	82.0	-1.9	-1.1
Atkinson index A(1) (%)	68.1	67.0	65.7	65.9	64.2	-5.7	-4.1
Atkinson index A(0.5) (%)	43.5	43.0	42.4	42.8	41.0	-5.7	-4.6
Regional Gini indices (%)							
Mature economies	38.2	38.9	39.1	38.8	41.9	9.7	7.9
China	32.0	35.5	38.5	41.8	42.7	33.5	20.6
India	31.1	30.1	31.4	32.4	33.1	6.3	9.9
Other Asia	44.5	44.3	46.6	41.8	45.0	1.1	1.6
M. East & N. Africa	41.8	42.0	43.5	39.4			
Sub-Saharan Africa		53.5	52.1	56.5	58.3		9.0
L. America & Caribbean	52.7	54.6	56.5	55.7	52.8	0.3	-3.3
Russia, C. Asia, SE Europe		48.3	40.1	41.8	41.9		-13.3
Decomposition by country: between-country contribution in % (change is in pp)							
GE(0) between contribution	83.2	80.1	78.2	77.9	76.7	-6.5	-3.4
Average annual incomes per capita (in 2005 PPP-adjusted USD), by percentiles							
Bottom 10%	201	203	217	228	251	24.9	23.3
P40-P50	552	620	715	766	941	70.6	51.8
P50-P60	791	877	975	1045	1359	71.7	54.8
P60-P70	1323	1353	1538	1616	2089	57.9	54.5
P80-P90	7414	7158	7177	7097	7754	4.6	8.3
P90-P95	12960	13150	13472	14221	15113	16.6	14.9
P95-P99	21161	21452	22660	24474	26844	26.9	25.1
Top 1%	38964	39601	46583	51641	64213	64.8	62.1
Average annual incomes per capita (in 2005 PPP-adjusted USD), by region							
World	3295	3287	3471	3631	4097	24.3	24.6
Mature economies	11457	12272	13366	15019	15832	38.2	29.0
China	484	572	789	1018	1592	228.9	178.3
India	538	560	638	642	723	34.4	29.1
Other Asia	671	804	882	943	1129	68.3	40.4
M. East & N. Africa	1773	1875	1974	1762			
Sub-Saharan Africa		742	719	779	762		2.7
L. America & Caribbean	3153	2982	3188	3024	3901	23.7	30.8
Russia, C. Asia, SE Europe		2757	2298	2544	4464		61.9
<i>Notes: For the decomposition by country, changes are in percentage points. For all other rows, changes are measured in percent (not annualised). Observations are weighted using population. The missing cells are deleted because of poor GDP/population coverage in particular benchmark years.</i>							

We could easily derive bootstrapped standard errors for the Gini index in order to account for sampling uncertainty (i.e. the fact that we have used a sample rather than the population). However, as Anand and Segal (2008) argue, these standard errors would not be appropriate, because they assume that there exists a single global household survey with a clearly defined sampling uncertainty. In contrast, we have combined a large number of national household surveys each of which has its own sampling uncertainty. As a result, plausible standard errors should probably be substantially bigger than the bootstrapped standard errors, making the observed changes insignificant.

As shown by Appendix Table 1, our estimates of the global Gini index are substantially greater than previous estimates in the literature.⁵³ The studies listed there differ fundamentally in their methodology, such as the use of national accounts aggregates, the type of PPP exchange rates and the interpolation for missing years. Most of the difference, however, is due to these studies' using the "old" 1993-based PPP exchange rates which give substantially lower price levels for China, India, Indonesia, Bangladesh and several other Asian countries, and hence imply higher incomes in those relatively poor countries. The closest study to ours is Milanovic (2012) who also uses surveys only, and applies the 2005 PPP exchange rates. Our estimate of the global Gini coefficient is greater than Milanovic's (2012), although the gap is falling over time, from 4.35 Gini points in 1988 to 1.76 Gini points in 2003. The direction of change between benchmark years is similar with the exception of the period between 1988 and 1993.

3.3 Alternative measures of inequality

We test for the robustness of these conclusions to different measures of inequality, such as the Generalised Entropy and Atkinson measures. The Gini index attaches a particular weight to inequality at different points along the income distribution. The Theil-L (or GE(0), or mean log deviation) index is particularly sensitive to differences in shares amongst low incomes, whereas the GE(2) index is sensitive to differentials at the top of the distribution (Cowell, 2009) and also sensitive to extreme values (Cowell and Flachaire, 2007). The Theil-T (or GE(1)) index is an intermediate case.

⁵³ Appendix Table 1 draws on Milanovic (2002, 2005, 2012), Bourguignon (2012) and Anand and Segal (2008). Note that Anand and Segal (2008) erroneously refer to Milanovic (2005) as GE(0), when in fact it is GE(1). For Milanovic (2002, 2005), we use the full sample, because this is closest to our approach. The decomposition is only available for the sample of countries that is common across all the years.

According to the top-sensitive GE(2) index, inequality increased between all benchmark years between 1988 and 2003. On the other hand, the bottom-sensitive Theil-L index shows falling inequality between 1988 and 1998, and a marginal, but probably insignificant, increase from 1998 to 2003. This appears to suggest that between 1988 and 2003, inequality amongst lower incomes was falling whereas it increased amongst higher incomes. Between 2003 and 2008, there has been a fall across the board, but a stronger change for the bottom-sensitive GE(0) measure.

We have computed the Atkinson (1970) index for three levels of inequality aversion ε . The higher ε , the stronger is the aversion to inequality in the distribution of incomes and the higher the weight attached to lower incomes. For $\varepsilon = 0$, society is indifferent to the degree of income inequality. With $\varepsilon = \infty$, only the position of the poorest group matters.

According to all three levels of ε , inequality is highest in 1988.⁵⁴ A(1) and A(0.5) agree on the relative rankings of the benchmark years (from lowest to highest inequality: 2008, 1998, 2003, 1993, 1988). For A(2), which is the highest level of inequality aversion considered here, 2008 has a higher level of inequality than 1998 and the level is not different between 2003 and 2008 (at least to one decimal point). Furthermore, $\varepsilon > 2$ would show increasing inequality between 2003 and 2008 (in contrast to all the other measures reported here). In sum, between 2003 and 2008, low incomes did not improve much leading to the same value of A(2) in 2008 and 2003, whereas the less inequality-averse A(1) and A(0.5) show an improvement.

3.4 Regional inequality and between-country decomposition of global inequality

The Gini index calculated across all individuals living in a region is highest in Latin America and Sub-Saharan Africa. The mature economies have seen a strong increase in the last benchmark year. Inequality in China has risen strongly between 1988 and 2008, by more than 10 Gini points. The increase in India has been much more moderate. The Gini index for Sub-Saharan Africa

⁵⁴ The level of inequality aversion can be interpreted as follows (Atkinson, 1975). Consider two people which are identical except for one having twice the other's income. Consider a transfer which takes away \$1 from the rich and gives a proportion x to the poorer person. An inequality aversion of 2 implies that we would accept a transfer in which only \$0.25 reach the poorer person for every \$1 taken from the rich. For $\varepsilon = 1$, this corresponds to \$0.50 and for $\varepsilon = 0.5$, we would require \$0.71.

increased by approximately 5 Gini points between 1993 and 2008. Within Middle East and Russia/C. Asia/SE Europe, inequality appears to have fallen. Inequalities within Latin America and Other Asia have remained virtually unchanged with some ups and downs in the intervening period.

We present a decomposition of the Generalised Entropy class measures, which are, in contrast to the Atkinson and Gini indices, additively decomposable.⁵⁵ We concentrate on the GE(0) index, because interpreting the within-group component as the residual inequality after equalising average incomes across countries, is only correct for this index out of the GE-class (Anand and Segal, 2008).⁵⁶ The between-country contribution has declined over this 20 year period, suggesting that countries, weighted by their populations, have become more similar.⁵⁷ In 2008, equalising mean incomes between countries while keeping the within-country distributions unchanged, would reduce global inequality by approximately 77%. Alternatively, equalizing all incomes within each country would reduce global inequality by 23% only. In other words, despite its relative decline, the between-county component still remains by far the more important source of global inequality.

3.5 Growth incidence curves

The bottom part of Table 3 displays growth in average incomes by income fractile. The group which has grown fastest is the one between the 50th and 60th percentile (growth rate of 71.7% over 20 years), followed by the P40-P50 group (70.6%) and the global Top 1% (64.8%). Perhaps a more useful way to illustrate this pattern is through a variant of the global Growth Incidence

⁵⁵ The Atkinson index is decomposable by population subgroups, whereas for example the Gini index is not. However, the Atkinson index is not additively decomposable in the sense that it can be broken up into a weighted average of the within- and between-group inequalities (Bourguignon, 1979).

⁵⁶ The GE(0) decomposition is in terms of population shares, whereas the decomposition of GE(1) uses income shares. Redistributing income among countries in order to equalise average incomes, would change income (but not population) shares. In that sense, only the interpretation of GE(0) is consistent because full elimination of one source of inequality (between- or within-inequality) will not affect the level of another.

⁵⁷ In an earlier draft, we also decomposed inequality by regions. The between-region contribution declined faster than the between-country contribution. This suggests that regions have become even more similar to each other than countries.

Curve (GIC) (Ravallion and Chen, 2003).⁵⁸ It compares the mean income of a given fractile group (e.g. the bottom 10%, the top 1%) in (say) 2008 with the mean income of the same fractile group in 1988. This is shown in Figures 1(a), 1(b) and 1(c) where the y-coordinate is simply the total growth rate between these two dates. A downward- (upward-) sloping GIC implies that economic growth has an equalising (desequalizing) effect on the income distribution, i.e. it is pro-poor (pro-rich). These are anonymous GICs because they wholly ignore the composition of people that find themselves in the same fractile group of the income distribution in two different years.

Figure 1 (a) shows the global GIC for the period 1988 to 2008. As we already saw from Table 3, growth was highest in the P50-P60 range. From around the 75th percentile, growth is lower than the growth in the global average. Then, for the top 1% of the global distribution, growth reverts to being higher than the average. This gives the GIC curve a distinct supine *S* shape, with two peaks, around the median and at the very top, and a trough around the 80-85th percentile. Because the GIC is everywhere above zero, the 2008 global distribution first-order stochastically dominates the 1988 distribution.

Figure 1 (b) repeats the global GIC for the separate 5-year periods between benchmark years. The GIC for 2003-2008 lies almost uniformly above the other periods suggesting that growth has been highest over this period. During 1988-1993 incomes declined particularly for the percentiles between the 70th and around the 88th. The quinquennial curves suggest that the supine *S* shape was present throughout the twenty-year period. The gains for the median and the top have been particularly strong in the last 2003-08 period whereas the losses for the groups around the 80th percentile have been exceptionally high in the first (1988-93) period.

The last part of Table 3 shows the growth in average income for the different regions of the world. Not surprisingly, China is the region with the strongest growth, average incomes tripling

⁵⁸ The original GIC, as defined by Ravallion and Chen (2003), shows the growth rate in incomes for the same percentile (e.g. the 10th percentile of the global distribution) between the initial and final period. In contrast, we compare the mean income of the same fractile group (e.g. the bottom 10%) over time. The rationale is an important one when there is a distributional change within a given fractile. To be very precise: The original GIC evaluated at 100 percentiles fails to show the growth in incomes if all or most of income gains are concentrated within the top 1%. Because the highest income group considered is the 99th percentile, the Ravallion-Chen GIC will show zero or almost zero growth throughout. Most of the growth in effect would not be recorded.

between 1988 and 2008. It is followed by Russia/Central Asia/SE Europe (only 15-year growth rates) and other Asia. The mature economies and India have grown at a very similar rate, in both cases superior to the world average. Latin America has grown at a (marginally) lower rate than the global average. Sub-Saharan Africa almost did not grow at all between 1993 and 2008. The regional ranking of growth thus clearly illustrates the success of China and the rest of Asia, a good performance of mature economies and India, and a very disappointing outcome for Africa.

Figure 1 (c) shows the 20-year GICs for five regions.⁵⁹ With the possible exception of the top 5% in Latin America, the GICs are everywhere above zero, so the 2008 distributions first-order stochastically dominate the 1988 distributions. Growth appears strongly pro-rich in China and less so in the mature economies and India, whereas the GIC is flat for Other Asia and displays no clear direction for Latin America.

While the global GIC showed relatively large gains for the portion of the distribution around the median, we need to recall that these gains were measured in relative (percentage) terms. But precisely because global income inequality is extremely high, and incomes at the top are several orders of magnitude greater than incomes at the median (in 1988, the average per capita income of the top 1% was close to \$PPP 39,000 while the median income was approximately \$PPP 600), the absolute gains are much greater for higher percentiles. Figure 1 (d) shows that the average per capita income for the top 1% increased by \$PPP 25,000 between 1988 and 2008, while the absolute gain at the global median was only \$PPP 400. The absolute gains among the poorer percentiles were even less. The overall outcome was thus that 44% of the increase of global income between 1988 and 2008 went to the top 5% of world population.⁶⁰

⁵⁹ The GICs by region are evaluated at decile groups (mean-on-mean) with the top decile being split into two ventiles. This is because for China and India, which are regions by themselves, we have at most 20 observations.

⁶⁰ These are, of course, not necessarily the same country-deciles (nor people) who were in the top 5% in 1988. We return to this issue in Section 5 where we discuss (quasi) non-anonymous growth incidence curves.

Figure 1(a): Global growth incidence curve, 1988-2008

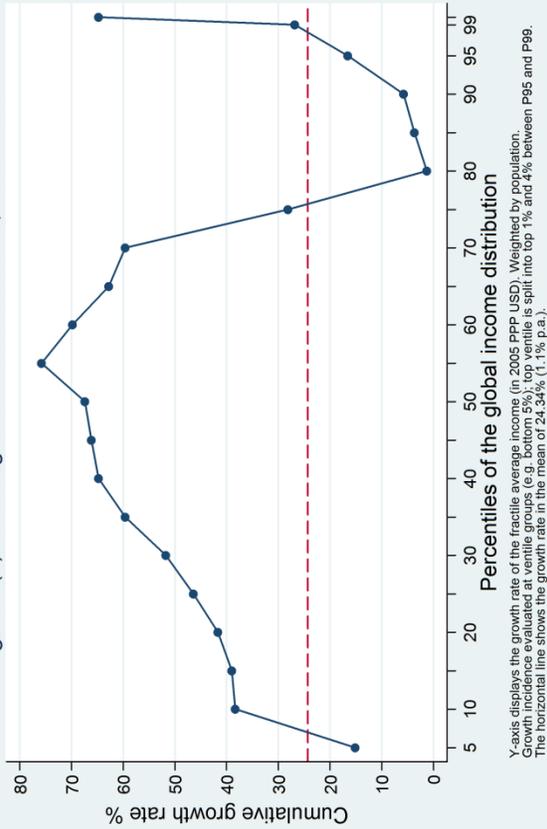


Figure 1(b): Global growth incidence curve, over time

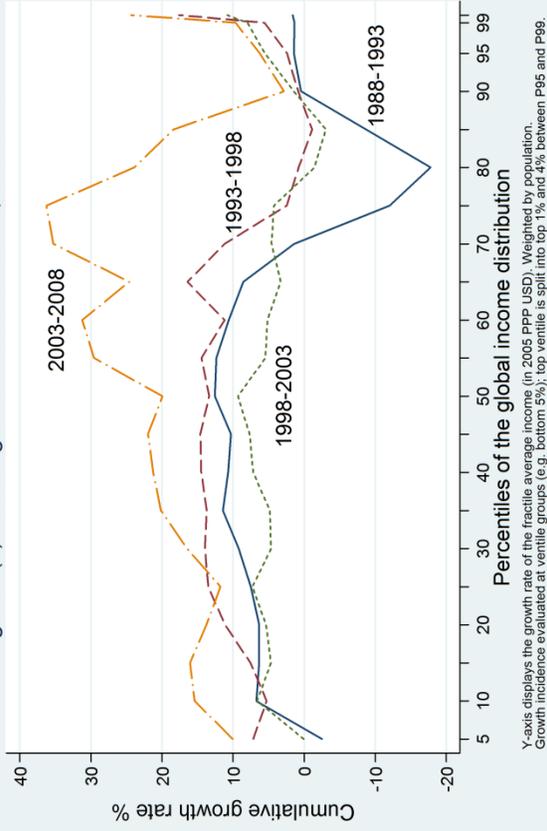


Figure 1(c): Global growth incidence curve 1988-2008, by region
Mature economies, Asia, and Latin America/Caribbean

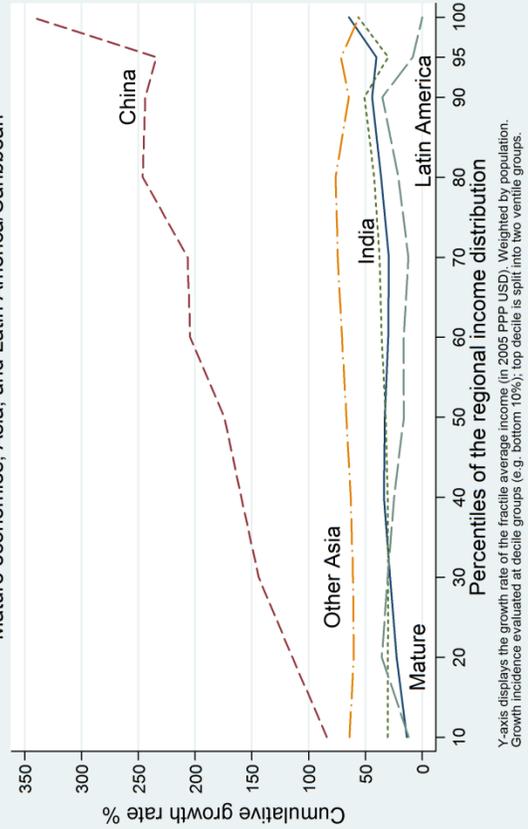
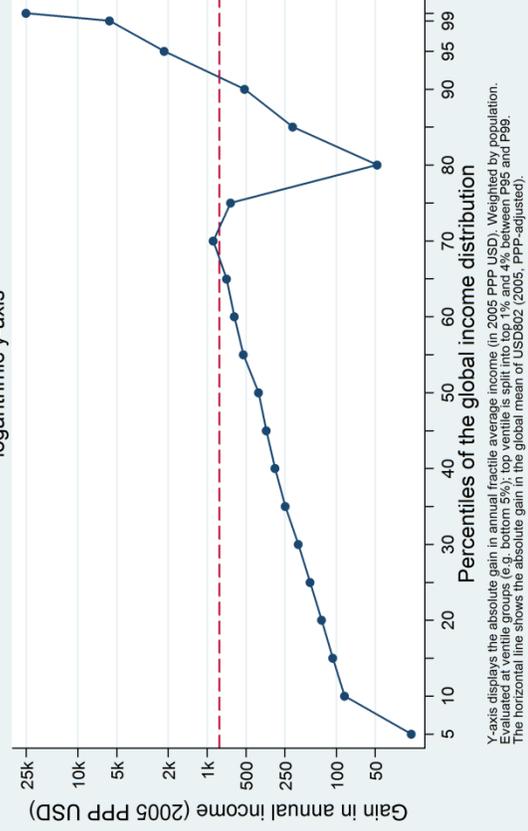
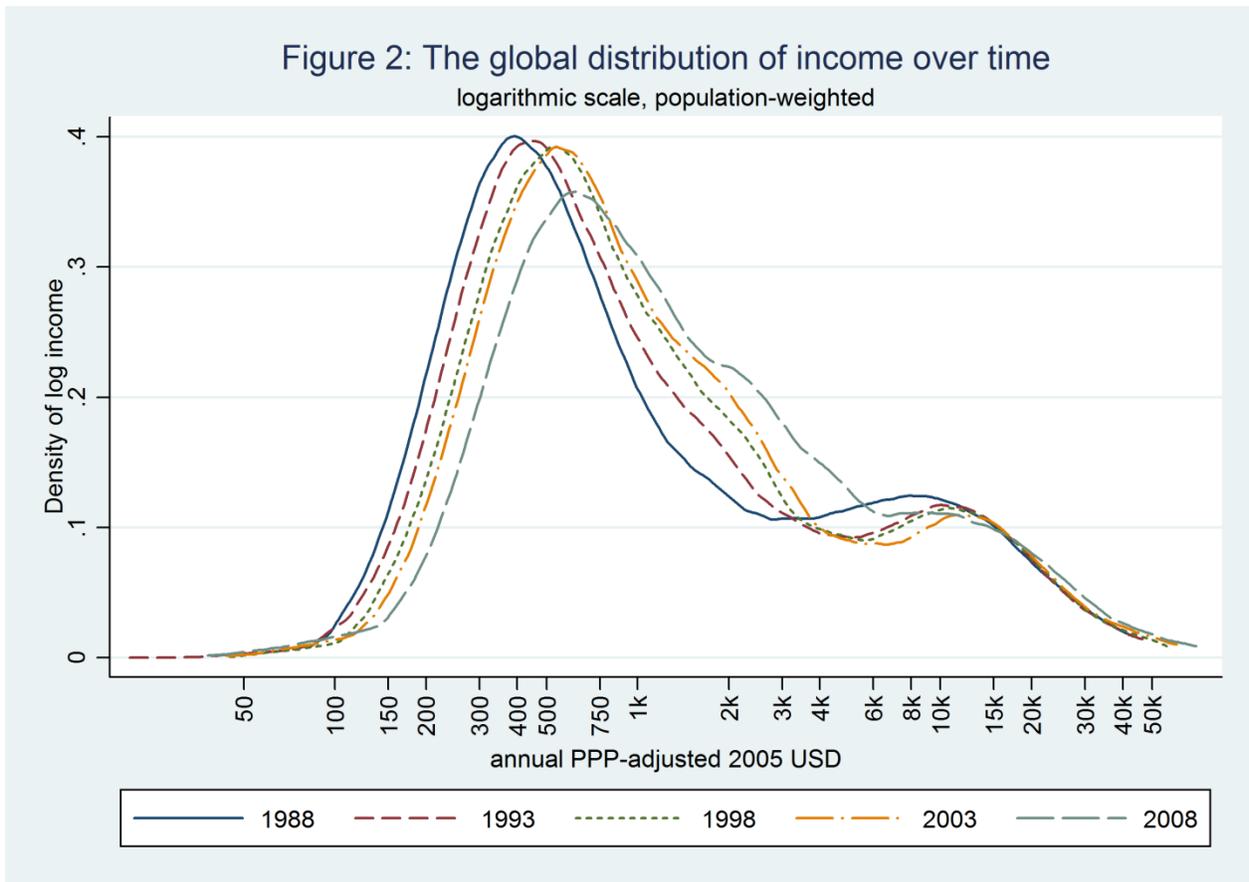


Figure 1(d): Absolute income gains, 1988-2008
logarithmic y-axis



3.6 The regional composition of the global distribution

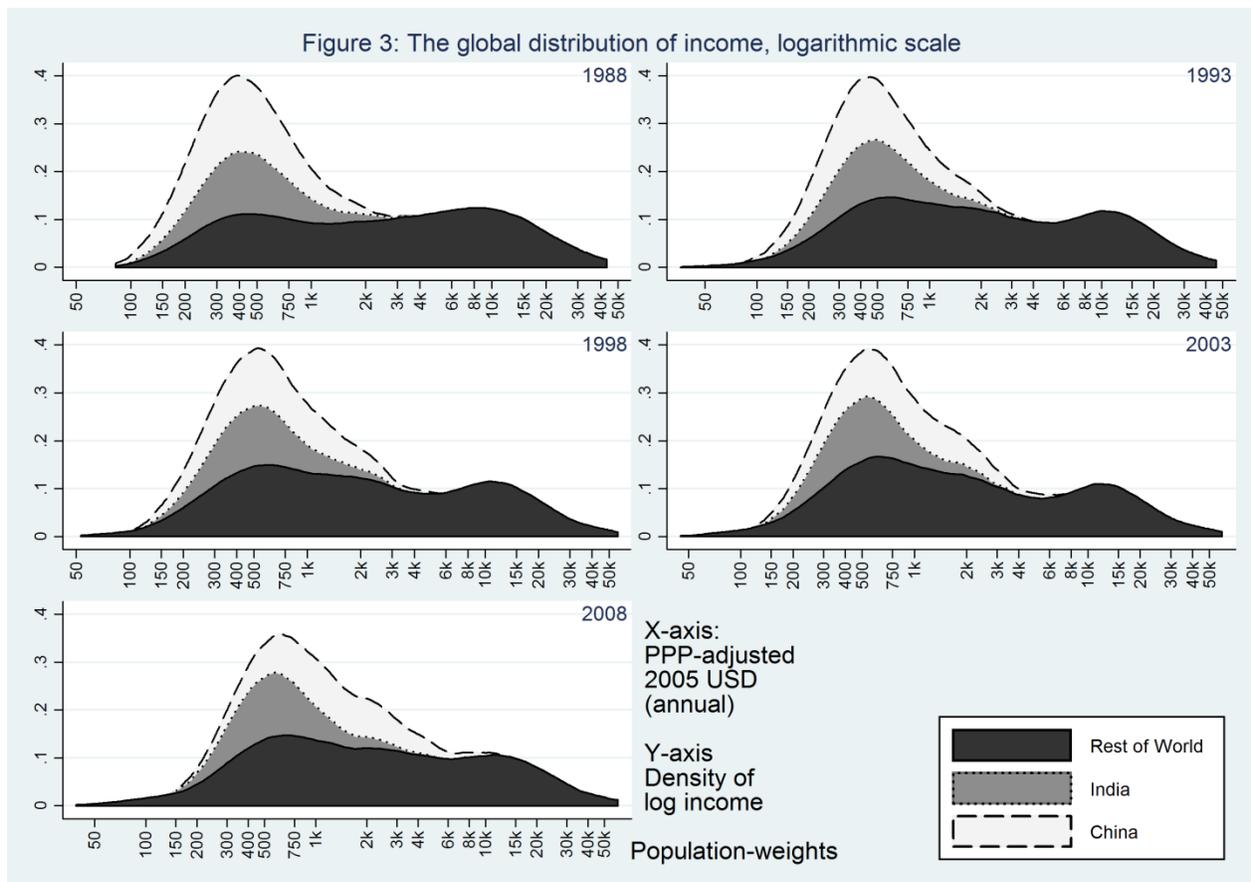
Figure 2 shows how the global distribution of income has changed over time.⁶¹ Income growth is shown by the rightward movement of the income distribution. The 1988 distribution appears to have two peaks, one around \$PPP 400 and another around \$PPP 10,000. In 2008, the second peak has disappeared and there is more mass around the \$PPP 3,000 mark. As implied by the almost universally positive five-year period GICs (Figure 1 (b)), the global distribution charts a rightward movement in every individual five-year period with the most striking development being the expansion of the proportion of the global population with incomes between \$PPP 750 and \$PPP 6,000 (that is, between approximately \$PPP 2 and \$PPP 16 per day). That population has expanded from 1.16 billion people or 23% of world population in 1988 to almost 2.7 billion or 40% of world population 20 years later.⁶²



⁶¹ We are using an Epanechnikov kernel and the default bandwidth, selected optimally by Stata.

⁶² The total numbers are for the entire world population, not only population covered by surveys here.

In order to disentangle these changes further, Figure 3 shows stacked kernel densities by regions.⁶³ Not surprisingly, the growth in China has had a profound effect on the global distribution. The change in the overall shape of the distribution appears to be driven by the upward income movement of the upper deciles in China. Both China and India have moved up along the income distribution while Sub-Saharan Africa (not shown in the figure) seems caught at the bottom.



The upward movement of China, because of its magnitude in terms of population and amount of growth, is particularly well illustrated in the stacked kernel densities. In 1988, Chinese population

⁶³ These charts have been created as overlaid (cumulative) kernel densities. Because the last density is shown on top we proceed in reverse order: The first density to be plotted is the global density including all regions (which is the same as Figure 2). Second, we plot the density for all regions, except China. We proceed by removing one region at a time. The area under the global density is 1. The other incomplete densities are scaled down according to the regional population share in a particular year. For instance, in 2008 the second density we plot is scaled down to $0.7828 = 1 - x$, where x is China's population share in 2008. We are using an Epanechnikov kernel and the default bandwidth, selected optimally by Stata. The bandwidth is allowed to vary for different years, but for every benchmark year it is the same across the different cumulative kernel densities.

was symmetrically distributed atop of the mode of the global distribution (exclusive of China). In other words, China and the rest of the world had about the same modal income. With each successive five-year period, Chinese distribution shifted more to the right (towards higher income levels) so much that by 2008 about four-fifths of the Chinese population has an income greater than the modal non-China global income. The income mode in China is now clearly greater than in the rest of the world. It is this rightward movement of the Chinese distribution that has most contributed to the change from a two-peaked global distribution in 1988 to a single peak distribution twenty years later. This has largely happened because China has “filled up” the relatively hollow part of the global income distribution between \$PPP 2,000 and \$PPP 6,000.

Figure 4 focusses on the change in the regional composition of the global income distribution between 1988 and 2008. The chart shows the regional composition of the population in each ventile of the global distribution.⁶⁴ As before, we can see a clear upward movement of China. The top decile in China reaches as far as the 17th ventile (i.e. between 80th and 85th percentile) of the global distribution in 2008, whereas in 1988, the richest Chinese were only between the 65th and 70th percentile. Conversely, in 2008 China has entirely graduated from the bottom 5% of the world, while in 1988 it made up almost 40% of the population in that group.

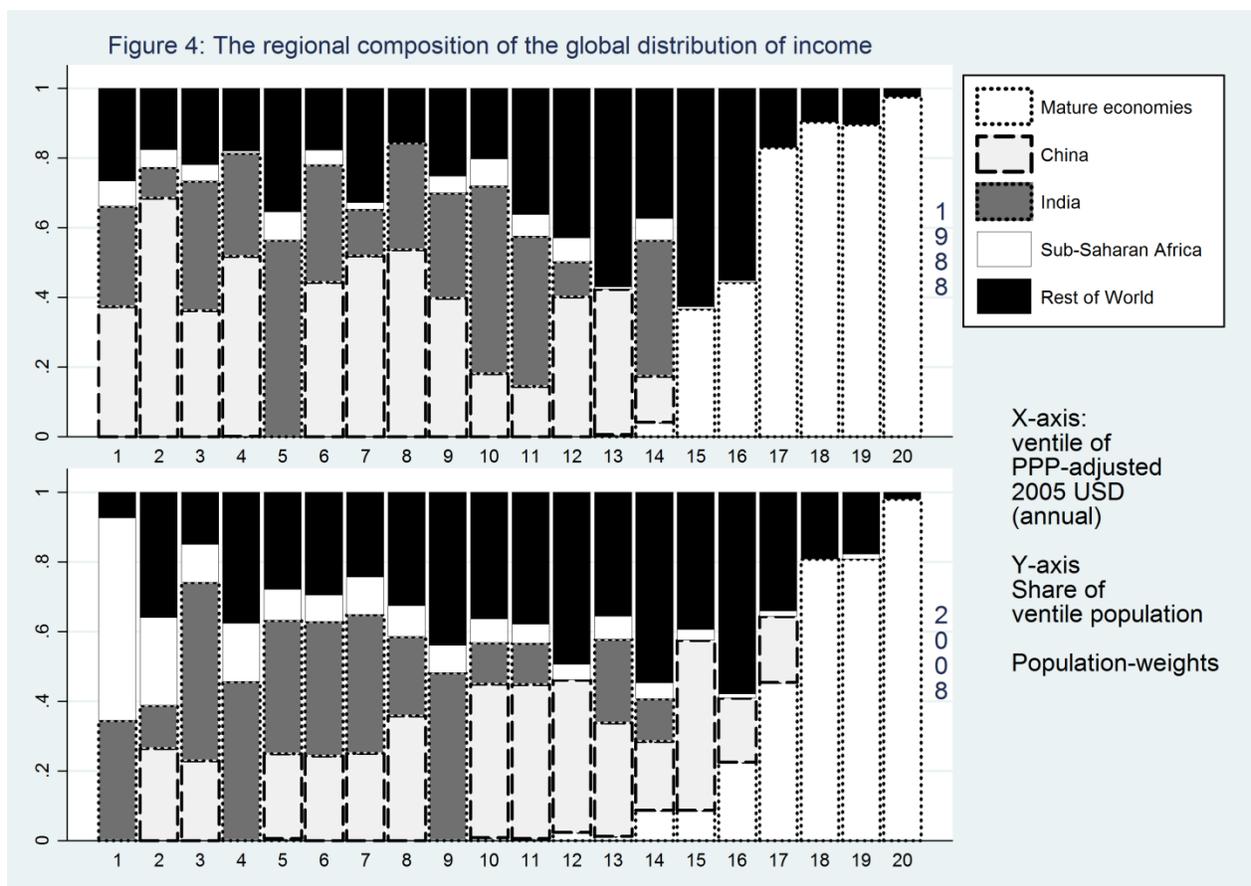
As the bottom incomes in China have moved up the global distribution, Sub-Saharan Africa and to a smaller extent India have expanded their population shares in the bottom ventile. The distribution of Sub-Saharan Africa is very spread out with some decile groups (from South Africa and Seychelles) reaching the top 10% of the global distribution. India has not moved dramatically which is explained by the fact that its growth rate has been similar to the global average.

The global 20-year GIC showed that fractile groups between the 75th and approximately 95th percentile grew slower than the global average (Figure 1(a)). In 1988, the percentiles between the 70th and the 85th (ventile groups 15, 16 and 17) originated primarily from the mature economies and Latin America, and to a smaller extent from the Middle East and North Africa. By 2008, China and to a lesser extent Russia⁶⁵ had moved into these percentiles, reducing the shares of the

⁶⁴ The ventile categories correspond to different absolute money amounts in 1988 and 2008.

⁶⁵ The first observation that we have for Russia is for benchmark year 1993.

mature economies and Latin America and the Middle East almost dropping out completely (not all regions shown separately). It is these compositional changes which explain the shape of the global GIC. The GIC does not track a particular fractile group, but rather compares the incomes of a given fractile in the *different* initial and final distributions. When comparing the top Chinese incomes in 2008 with the Latin American incomes in 1988, we obtain a below-average growth rate. This is despite the fact that the top Chinese incomes have grown substantially faster than the global average. However, this topic, the (quasi) non-anonymous Growth Incidence Curves, which keep the composition of fractiles the same as in the original year, is discussed in Section 5 below.



4 Accounting for missing top incomes in the computation of global inequality

So far in the analysis we used our main cross-sectional sample and only the information contained in the surveys, i.e. we did not adjust survey means to national accounts. In this part of the analysis, we test the robustness of our conclusions to (1) anchoring to national accounts, that is distributing the “excess” income from the gap between national accounts consumption and the household survey mean either across the entire distribution or to the top decile only, and (2) using a Pareto interpolation to “elongate” the distribution of the top decile.

	Benchmark years					1988-2008	1993-2008
	1988	1993	1998	2003	2008	change (pp)	change (pp)
(1) Survey data only	72.5	71.8	71.9	71.9	69.6	-2.9	-2.2
(2) Private consumption instead of survey mean	71.5	70.5	70.6	70.7	67.6	-3.9	-2.8
(3) Private consumption with Pareto imputation	71.8	70.8	71.0	71.1	68.0	-3.7	-2.8
(4) Private consumption with top heavy Pareto imputation 1/	76.3	76.1	77.2	78.1	75.9	-0.5	-0.2
Number of surveys	63	105	112	126	114		

Notes: Observations are weighted using population. All calculations are done across the sample of 520 country-years for which private consumption from national accounts is available.

1/ The entire gap is allocated to the top 10%.

Table 4 presents our results. Since here we replace survey means with private consumption from national accounts (only when the latter is greater, otherwise we keep survey means⁶⁶) we lose country-years for which we do not have national accounts information. This leaves us with 520 surveys across the five benchmark years instead of 565. In addition, we now treat China, India and Indonesia as single countries because we do not have separate national accounts information for rural and urban areas. The new baseline Gini for this new sample (row 1 in Table 4) is very similar to the Gini for the full sample (Table 3): the difference is less than 0.5 Gini points except

⁶⁶ This means that we use the maximum of per capita private consumption and the survey mean. In other words, when the survey mean is greater than private consumption we keep the survey mean. Atkinson (2007) replaces survey means by national account aggregates in all cases. We adopt a different approach because we focus on the underestimation of top incomes in household surveys. In other words, we assume that surveys cannot overestimate, but only underestimate, overall income of households.

in 2008, when the new baseline Gini is 0.9 Gini points lower than the one obtained from the full sample.

First, we replace survey means by private consumption from national accounts, a “re-anchoring” that has been often done in the literature except that it was typically done using GDP rather than private consumption. Such “re-anchoring” obviously leaves within-country inequality unchanged. The entire effect on global inequality comes from the change in the between component (and, in the case of the Gini index, indirectly from the change in the overlap component). The between component obviously changes because the country-means change. *A priori* the direction in which we expect the global Gini to change is unclear. However, previous calculations have mostly shown that re-anchoring to GDP (with no other adjustments) tends in more recent years to lower global inequality (see Milanovic, 2005, p. 118, and Appendix Table 1). We find here the same result. As can be seen in Table 4 (row 2), Gini declines by approximately 1 point except (again) in 2008 when it goes down by 2 Gini points. Intuitively, the reason for the downward change is that survey underestimation is greater in poorer (population-weighted) countries. In other words, poor countries appear less poor when we replace surveys with private consumption from national accounts.

Second we assume that the distribution of the top 10% can be approximated by a Pareto distribution, which is similar to Atkinson (2007) approach.⁶⁷ This “proportional adjustment with Pareto tail” allows within-country inequality to increase, while between-national inequality remains the same (since means are unchanged), and the overlap term remains the same or increases. The overlap component is likely to increase because the “elongation” of the distributions will tend to increase the “mixing up” of incomes of people from poorer and richer countries (or in the extreme case, leave it unchanged).⁶⁸ Consequently, we would expect the overall Gini, compared to the one from row 2, to be greater. This is indeed the case (see row 3) although the increase is quite moderate: at most ½ Gini point.

⁶⁷ Our approach differs from Atkinson’s in three respects: First, we use private consumption, whereas Atkinson uses GDP. Second, we keep survey means when they are greater than national accounts means, whereas Atkinson does not. Third, Atkinson assumes the Pareto distribution only for the top 5%.

⁶⁸ The overlap component remains unchanged if the country-means are very far apart and the Pareto imputation only has a small effect. For example, if we have only Switzerland and Congo in the sample, the Pareto elongation of the top of the Congolese distribution may not yet place any Congolese fractiles within the range of the Swiss distribution (that is, all Swiss fractiles will still have higher incomes).

Finally, we apply our “top heavy adjustment with Pareto tail”, where within-country inequalities are allowed to increase even further. This is done by allocating the entire gap between private consumption and the survey mean to the top decile and applying (as in the previous case) a Pareto “elongation”. It should be intuitively clear that by increasing the share of the top decile, we make income inequality greater and the Pareto constant lower.⁶⁹ In some cases, as we discuss in Appendix 2, such an adjustment may seem excessive an issue which we need to address more comprehensively in future work. For example, if the survey mean is equal to only 50% of private consumption (which is similar to the value observed in India) then simply “ascribing” these 50% to the top decile is probably excessive. Incomes of lower deciles are likely to have been underestimated as well.⁷⁰ This is why we consider “top heavy adjustment with Pareto tail” to be an extreme case. The results (row 4, Table 4) show that the Gini increase (compared to the “proportional adjustment with Pareto imputation”, row 3) is now between 4.5 and almost 8 points. Again, the most dramatic change occurs in 2008. In Appendix 2 we test for the robustness of these values to a more plausible range of the Pareto coefficients.

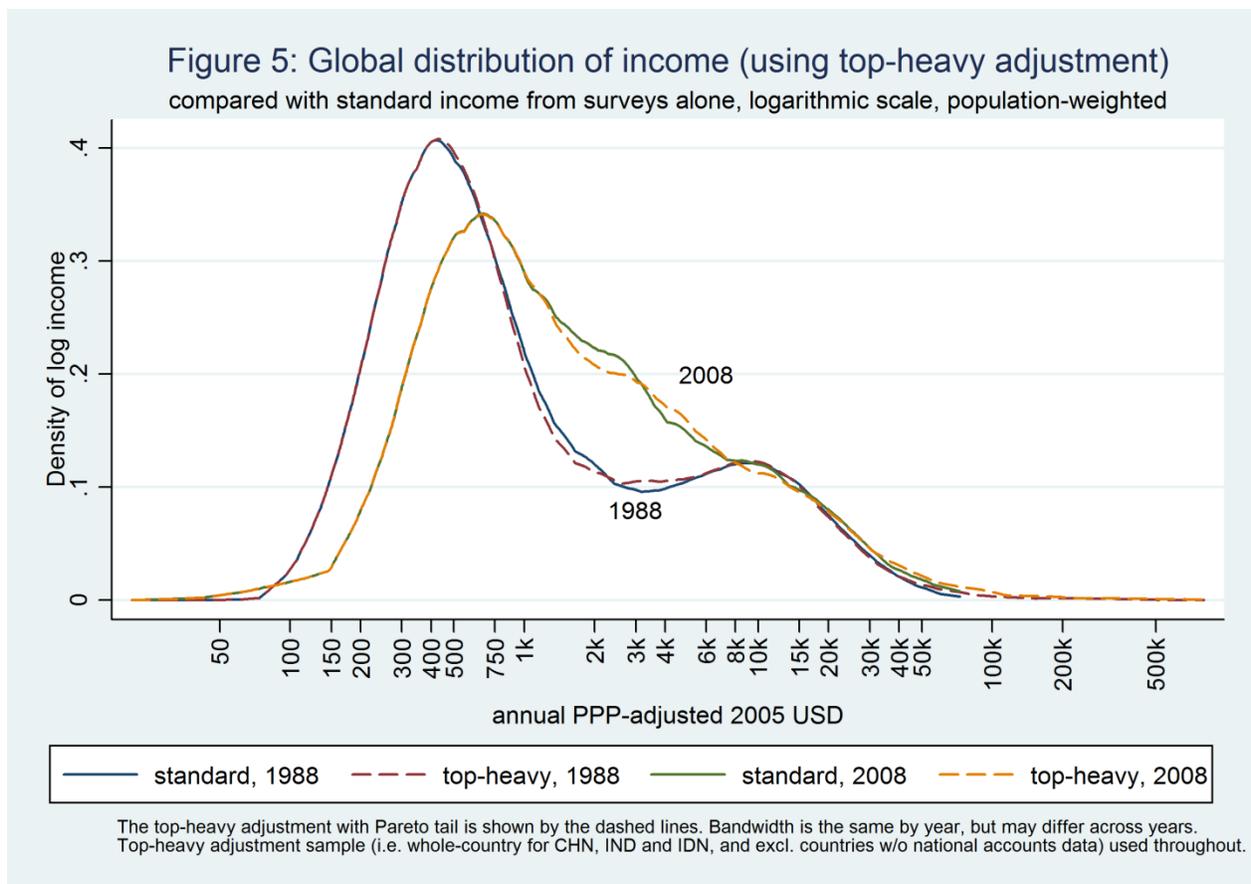
In summary, global Gini calculated from survey data alone is reduced by between 1 and 2 Gini points when we replace survey means by private consumption from national accounts (thus allocating the gap proportionally across national income distributions). When we also assume a Pareto upper tail, the overall Gini barely changes: it increases by about 0.5 Gini points.⁷¹ Only if we increase inequality further by not allocating the gap proportionally, but imputing it to the top decile only, does global Gini increases substantially, by between 4.5 and almost 8 points. This therefore, we believe, sets the range within which the “true” global Gini is likely to be. In 2008, for example, that range is between 68% and 76%. We tend to believe that it is closer to the upper bound but there is no way to prove it.

⁶⁹ A smaller value of the Pareto coefficient implies higher inequality.

⁷⁰ Of course, this is assuming away any definitional differences and other reasons why means in national accounts and household surveys might differ.

⁷¹ However, for other applications, such as the number of people around the world above a certain income level, the Pareto imputation might make a substantial difference (see Atkinson, 2007).

The main result of this exercise is not, however, the range of the level of global inequality, but the likely change between 1988 (or 1993) and 2008. As the last two columns in Table 4 make clear, with a “top heavy” adjustment the decrease in global inequality, present when we use all other adjustments, almost entirely dissipates. The change in the global Gini over these 20 (or 15) years is now merely -0.2 or -0.4 points. The reason is that over the period the gap between national accounts and survey means has risen from an average of 19 percent in 1988 to 25 percent in 2008⁷² (see Appendix Table 2). When we allocate this rising gap entirely to the top tail, we obtain an increasing within-country, and ultimately global, inequality. Before, we argued that the change in the global Gini index observed for the full sample (and using incomes directly reported in the surveys), was probably not robust to plausible standard errors. This robustness check further supports a more cautious view about the decline in global inequality: if indeed surveys tend to underreport incomes at the very top, it could well be that global inequality, measured by the Gini index, has not gone down during the twenty-year period considered here.



⁷² These are unweighted gaps. In India, the gap has increased by almost 25 percentage points (of national private consumption). But the situation is not much better elsewhere: the mature economies, Middle East, and Sub-Saharan Africa all display increasing gaps of around 10 percentage points (see Appendix Table 2).

Using the income reported in surveys we concluded above that the global income distribution had moved from a twin-peak to a single-peak. This also holds for the global distribution of income which adjusts for missing top incomes (Figure 5). Not surprisingly, the top-heavy Pareto imputation stretches out the top tail and makes it thicker from around \$PPP 40,000. Further down the distribution, mass appears to shift from around \$PPP 3,000 to \$PPP 5,000, which are the top fractiles in the poor countries.

5 Changes over time: Who are the winners and losers?

In this section we return to the issue of which country-deciles have contributed to the overall changes in the global distribution. The evidence on the changing regional composition (Figures 3 and 4) was a first step into this direction, although it did not say anything about the movement of individual country-deciles. This section tries to identify the particular country-deciles which have gained or lost most over this period.

The sample used in the main text includes all countries which are observed in 1988 and 2008, even if there are gaps in the intervening benchmark years. In the Appendix, we have replicated the results using all countries observed in 1993 and 2008, which improves the coverage particularly in Sub-Saharan Africa and Russia/Central Asia/South-East Europe.

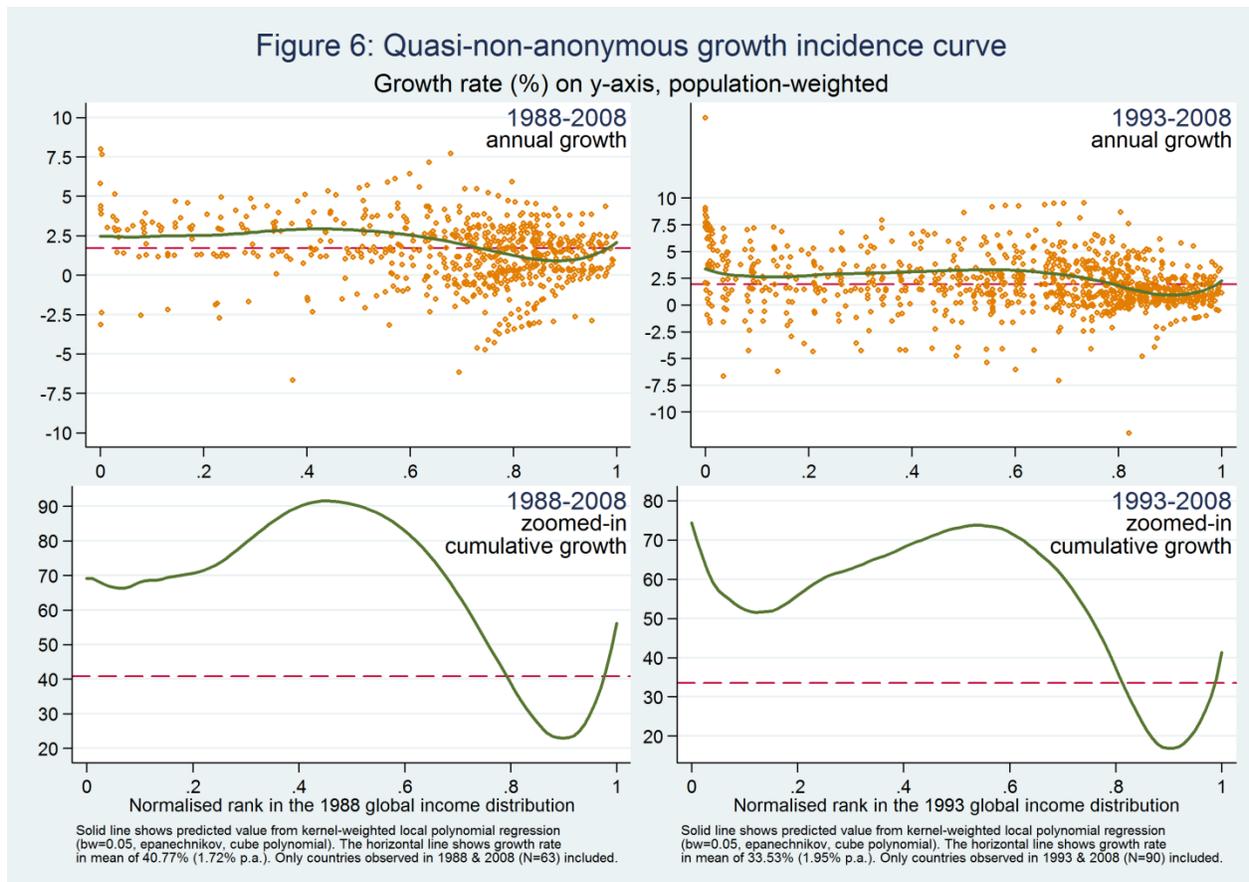
5.1 Quasi non-anonymous growth incidence curves

In order to identify the winners and losers, we first consider “income mobility profiles” (van Kerm, 2009), or “non-anonymous growth incidence curves” (Bourguignon, 2011a and Grimm, 2007). When applied to individual-record data, the distinction into anonymous and non-anonymous GICs is straightforward: The (standard) anonymous GIC compares the incomes of say the 20th percentile in the initial and final period distributions. As long as there is some mobility in the distribution, the individuals at this percentile might be different. By contrast, the non-anonymous growth incidence curve is the (non-parametric) regression of income growth against the rank (percentiles) in the initial distribution.⁷³ Because this growth rate is obtained for each individual, it is non-anonymous, taking into account the joint distribution of initial and final incomes. However, in our case the unit of analysis are income deciles of a particular country, so while we preserve the identity of a particular country-decile, these deciles are defined over different people. Hence we refer to our figures as “quasi-non-anonymous” growth incidence curves.

Figure 6 shows the quasi-non-anonymous GICs for 1988-2008 and 1993-2008. It plots the growth over the next twenty (fifteen) years against the normalised fractional rank in the 1988

⁷³ A special type of non-anonymous GIC would consist in showing the suitably weighted average growth rate of all individuals who have belonged to a given percentile in the initial year (and doing it, of course, for all 100 percentiles). Such a special non-anonymous GIC can consist of a comparison of weighted mean income of the same people (belonging to a given percentile) in the final and initial year.

(1993) cross-sectional global income distribution. In order to exploit both the cross-sectional and the panel dimensions of our data, each observation included in the quasi-non-anonymous GIC is ranked in the complete cross-sectional 1988 distribution (population-weighted) (not only amongst the 63 countries observed in 1988 and 2008).⁷⁴ For each time period, we present charts with and without the scatterplot.⁷⁵ The scatterplots show the wide dispersion of growth rates around the fitted line. Judging from the scale of the y-axis, the dispersion seems to have increased, but this might be driven by one outlier close to the bottom ranks in 1993-2008. The fitted line (a kernel-weighted local polynomial regression) is shown on a more detailed y-axis in the bottom panel for 1988-2008 and 1993-2008 periods.



⁷⁴ Fractional ranks are derived from a smooth cumulative distribution estimator which ensures that the mean rank is 0.5, estimated using the fracrank Stata routine by Philippe van Kerm.

⁷⁵ There are two reasons why observations in the scatterplot are not equally distributed along the horizontal axis, and appear concentrated amongst the upper ranks. First, the ranks are computed in the cross-sectional distribution, whereas the scatterplot includes only those country-deciles which are observed in 1988 and 2008, many of which are from rich countries. Second, country-deciles are weighted by population size in calculating the ranks, whereas this information does not show up in the scatter points.

It is immediately apparent that the shape of the quasi non-anonymous curves is very similar to the shape of the anonymous GIC (Figure 1 (a)): they all display a supine *S* shape driven essentially by very slow growth around the 80th and 90th percentile of the global income distribution and local maxima around the median of the income distribution and for the very top. However, if we compare the 1988-2008 results, it is clear that the gains among the country-deciles that were in the top 1% in 1988 were less than the gains we obtain by simply comparing income levels of the top 1% in 2008 and 1988. This is expected: not every country-decile that was in the top 1% in 1988 managed to have high growth over the next 20 years. Similarly, some country-deciles that were not among the top 1% in 1988 and have exhibited high growth are now (in 2008) in the top 1%. We find the equivalent result for the 1988 poorest country-decile whose growth rate was higher than what we found from the anonymous GIC. Furthermore, some of these differences might arise from restricting the sample to those countries which are present both in 1988 and 2008.

Over the 1988-2008 period, growth was the highest for those country-deciles around the 40th percentile of the 1988 global distribution, and lowest around the 85th percentile. Groups that were most successful come overwhelmingly from China and India; groups that were least successful are predominantly from the mature economies. Thus, ¾ of the population that was between the 36th and 45th global percentile (inclusive of the 45th) in 1988 belonged to the country-deciles, generally around the middle of national income distributions, from China and India. If we include other Asia too, the percentage of people belonging to this most successful group reaches 90%. Chinese deciles, for example, multiplied their incomes by a factor of 2.7-2.8.

In contrast, the country-deciles between the 81th and 90th (inclusive of the 90st) percentile in 1988 are overwhelmingly from mature economies, and come from the lower halves of their national income distributions. Out of total 420 million people belonging to this group, about 365 million are from the mature economies (or differently, 135 out of 165 country-deciles). Even when we exclude from the mature economies those that in 1988 were Communist, the share of the “traditional” rich economies among this group is still very large: 78 percent of people. Some examples with particularly low real growth rates among rich economies include almost the entire lower halves of the income distributions in Austria, Germany, Denmark, Greece and the United

States. They all had overall 20-year growth rates of less than 20% which translates, in the best case, as 0.9% per capita annually.

Over the 1993-2008 period, growth was the highest around the 60th percentile. The shape of the fitted line also seems to have changed with some high growth rates amongst the very bottom 1993 ranks. This is due mostly to the inclusion in 1993 of Russia (absent in 1988) whose low deciles experienced a period of high growth between 1993 and 2008.

5.2 Most successful and least successful

We next look directly at the 20 biggest gainers and losers (country-deciles) over this twenty-year period (Table 5).⁷⁶ We rank country-deciles according to the average of the (annualised) 5-year growth rates between 1988 and 2008.^{77 78}

Between 1988 and 2008, the top twenty country-deciles (but one) have all grown in excess of 5% annually, which means that their incomes have increased by at least 2.65 times. For the deciles that have grown at around 8% annually, real incomes have been multiplied by a factor greater than 4.5. Almost all urban Chinese decile groups are amongst the twenty fastest growing world deciles. Moreover, they are neatly ranked with highest deciles having grown the fastest. This illustrates the already observed disequalizing pattern of growth in China. Rural China has grown slower but it is interesting to observe that the top two rural deciles are also among the twenty winners, and again are ranked, with the rural top decile growing faster than the second highest rural decile. Overall, the remarkable character of Chinese growth - extremely fast and disequalizing - is well illustrated in these data: half or more of the most successful country-deciles come from China. Notice that this has nothing to do with population size because each county-decile is treated equally here. In effect, because of the huge size of Chinese deciles, one

⁷⁶ Country-deciles are numbered 1 to 10, with 1 being the bottom decile.

⁷⁷ This excludes any countries which have no 5-year intervals. Cyprus (for 1988-2008) and Niger (for 1993-2008) are observed in the initial and final years, but have no 5-year growth intervals in between.

⁷⁸ By contrast, Figure 6 shows the (annualised) growth in decile income between 1988 and 2008 on the y-axis, which says nothing about the permanency of that growth. For example, a decile might show a particularly high growth rate between 1988 and 2008 because it had an exceptionally large 2008 income (or, equivalently, was unlucky in 1988). Incidentally, the two rankings are very similar.

could argue that it would be harder to increase the average income of such large masses of people than would be the case in a smaller country. Other than China, the other fastest growing deciles are the relatively poor deciles in El Salvador, Costa Rica, Ireland, the UK and Chile.

Table 5: Winners and losers in terms of average annualised growth (1988-2008)			
20 biggest gainers (best at the top)		20 biggest losers (worst at the bottom)	
decile	growth	decile	growth
El Salvador-1	9.6%	Bulgaria-1	-4.4%
Costa Rica-1	7.9%	Lithuania-6	-4.5%
China-urban-10	7.7%	Romania-8	-4.6%
Ireland-1	7.2%	Lithuania-4	-4.7%
Mexico-10	6.6%	Estonia-2	-4.7%
China-urban-9	6.5%	Romania-7	-4.9%
China-urban-8	6.1%	Estonia-3	-4.9%
China-urban-7	5.9%	Lithuania-5	-4.9%
China-urban-6	5.7%	Estonia-7	-5.1%
China-rural-10	5.6%	Romania-6	-5.1%
Ireland-2	5.6%	Romania-5	-5.3%
China-urban-5	5.6%	Estonia-4	-5.3%
UK-1	5.5%	Romania-4	-5.5%
China-urban-4	5.4%	Estonia-6	-5.5%
Ireland-3	5.3%	Romania-3	-5.6%
China-rural-9	5.1%	Estonia-5	-5.6%
China-urban-3	5.1%	Romania-2	-5.8%
Ireland-4	5.1%	Bolivia-1	-6.0%
Chile-1	5.0%	Estonia-1	-6.4%
China-urban-2	4.9%	Romania-1	-7.4%

Notes: Only for countries observed in 1988 and 2008, and which have at least one 5-year growth interval. Deciles are numbered 1 to 10, with 1 being the bottom decile.

Among the bottom, there is a similar concentration. With the exception of Bolivia, East European deciles are almost solely filling up the ranks of the most unsuccessful groups.⁷⁹ Practically all deciles in Romania and Bulgaria seem to have experienced negative growth. An average decline of 3-4% per annum translates, after 20 years, in a real income loss of about one-half. The results

⁷⁹ Ranking by 20-year growth rates changes the results slightly as the bottom deciles of Honduras and Paraguay are amongst the twenty least successful deciles. This is also the case for the least successful deciles in terms of 15-year growth rates (Appendix Table 4), where Kenya and Tanzania have performed particularly poorly.

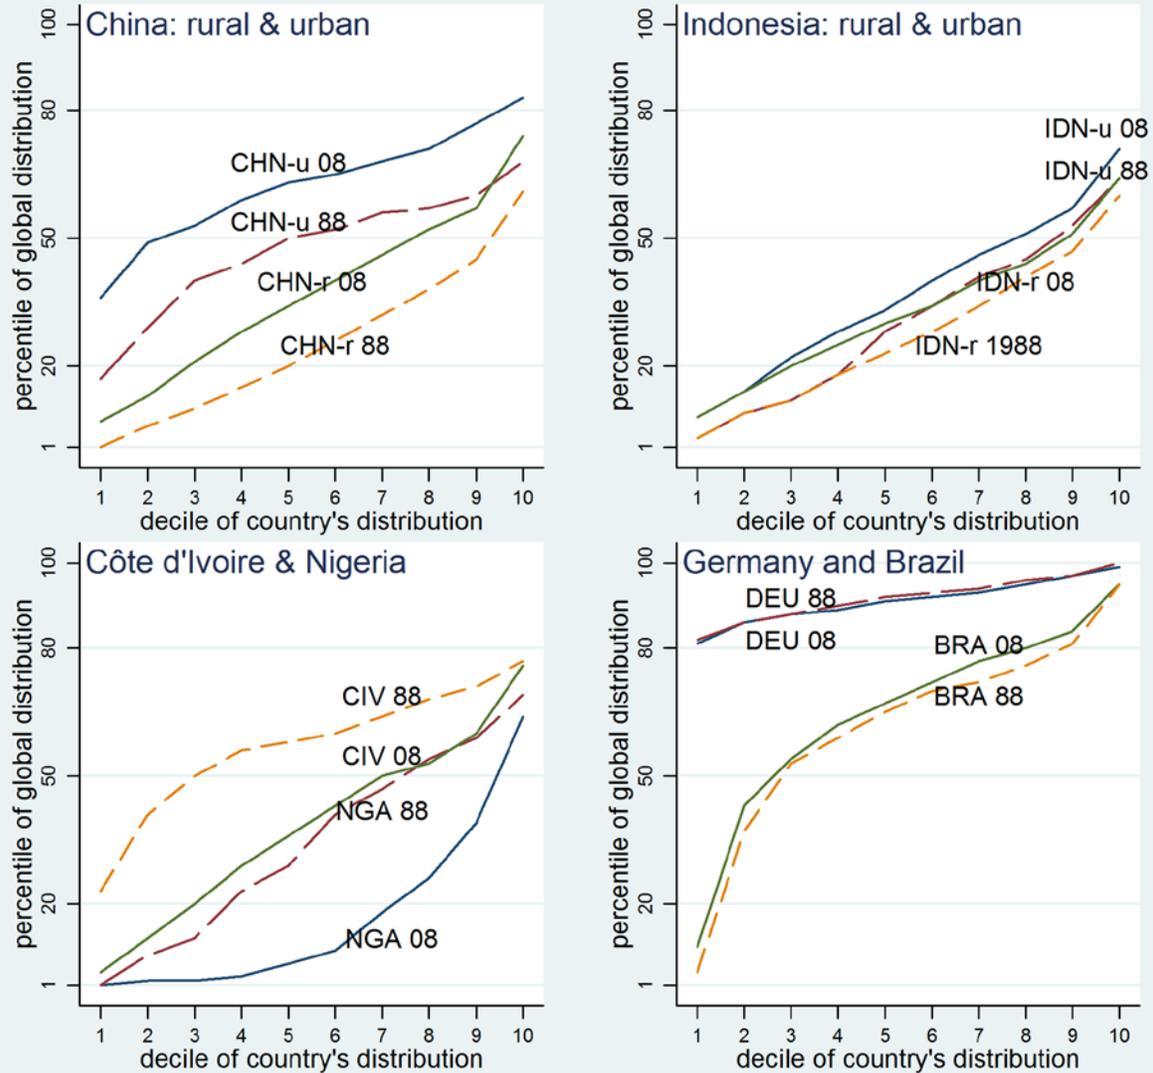
are different for the period 1993 to 2008 (Appendix Table 4).⁸⁰ The country-deciles that have lost most between 1993 and 2008 are mostly from Sub-Saharan Africa and Latin America, with East European deciles having recovered. These results are consistent with the disastrous income declines in Eastern Europe between 1988 (the beginning of the reform process) and mid-1990s followed by, in some countries sharp, and in other countries, rather slow recovery.

Another simple way to evaluate the success of various country-deciles is to compare their positions in global income distribution. This can obviously be done for every country-decile and every year. In Figure 7, we do it for several selected countries (2008 values are always drawn as a solid line, 1988 values as a dashed line). The top left panel illustrates the already discussed remarkable upward mobility (in the global income distribution) of Chinese rural and urban deciles. In 2008, the Chinese top urban decile is at the 83rd global percentile while twenty years earlier it was at the 68th. In other words, if somebody stayed in that same decile in China over 20 years and his income followed the average growth path of the decile, he would have leap-frogged, in terms of income, more than 900 million people worldwide. It is interesting that in 2008, the Chinese top *rural* percentile is at a higher global position than the Chinese *urban* percentile was in 1988. The same development, although less dramatic, is illustrated by the improvements in the position of Indonesian rural and urban deciles shown in the upper right panel of Figure 7.

But when we turn to the bottom left panel the situation is exactly the reverse. There we see that Nigeria's and Côte d'Ivoire's deciles have almost uniformly slid down the global income distribution. It is only the top Ivoirian decile that has managed to preserve its 1988 position. Finally, in the bottom right panel, we show the evolutions in Germany and Brazil. The position of German deciles has remained very high and displays very little change. Brazil is an example of a reasonably fast growth across its income spectrum, and improvements in the distribution so much so that all its deciles, except the highest, are now placed globally higher than they were in 1988.

⁸⁰ The bottom decile in Japan declined by 4.5% per year on average, which seems a very substantial decrease. Over the same period, average incomes in Japan (in our data) declined by 1.35% per annum, whereas GDP per capita (in 2005 \$PPP) grew by 0.88% per annum. Only in the most recent benchmark year we were able to obtain micro data from LIS for Japan. Before that the data are based on tabulations, which might explain the very substantial decline between 1988 and 2008 for the bottom decile.

Figure 7: Position of national deciles in global distribution (1988 & 2008)



In this context, China-rural, China-urban, Indonesia-rural and Indonesia-rural are treated as “national” income distributions. 2008 values are always drawn as a solid line, 1988 values as a dashed line.

6 Conclusion

The paper has provided evidence on the evolution of the global income distribution during a crucial period of accelerated globalization spanning the period from the end of Communist regimes in Eastern Europe and the Soviet Union to the beginning of the global financial crisis. In many respects, this might have been the most globalized (and most optimistic, in the sense of trust in the potentials of globalization) period ever in history. It is very important to study what were its effects on the level and distribution of income among the world population. Our results confirm earlier findings that the level of global inequality remains high, with a Gini of around 70%, and while inequality appears to have declined in the most recent years, these changes are probably not robust to plausible standard errors (if one could formulate and calculate them).

A robustness check whereby missing incomes, defined as the (positive) gap between per capita household consumption from National Accounts and mean per capita income from household surveys, are allocated to the top 10% of individual recipients, and top incomes are supposed to follow a Pareto distribution, returns estimates of the global Gini coefficient which are about 5 points higher. Most of the increase is due to the allocation of the entire gap to the top decile, not to the Pareto-elongation of the top of the income distribution. A value of 75% may be considered an upper bound on the “true” global Gini.

The approach whereby the entire gap is allocated to the top national deciles seems to us a more realistic one than the alternative used so far in the literature of spreading the gap evenly. In effect, we argue that the two problems that have recently been discussed separately - namely the large, and in some cases, growing gap between National Accounts consumption and means from household surveys, and the realization, due to the results obtained from fiscal data, that surveys tend to underestimate top incomes - are really one and the same problem.

Our approach returns another important result. Global inequality, measured by the Gini index, might not have gone down at all if the entire (or most) of missing income comes from the underestimation of incomes at the top of national income distributions. This important issue

obviously requires more research but it highlights the possible global effects of misreporting among the top that have been already noticed in individual countries.

The shape of the global income distribution has also changed during the 20 years considered here (referring to the baseline scenario without adjusting for missing top incomes). In 1988, the global income distribution displayed a twin-peak shape which has since disappeared mostly thanks to the high growth of China whose deciles have “filled up” the area between \$PPP 2,000 and \$PPP 6,000 that was relatively “hollow” in 1988. The period has also witnessed a remarkable increase in what may be called a “global median class”, with incomes ranging from \$PPP 2 per capita per day to \$PPP 16 per capita per day: the share of the global population belonging to that group has increased from some 23% to 40%.

Particularly important is the shape of the anonymous and quasi non-anonymous growth incidence curves for the period 1988-2008. They both show large relative gains around the median of the global income distribution that accrued mostly to the middle or upper-middle income deciles from Asia, and especially from China. By contrast, the lowest real income gains were registered in the area around the 80-85th global percentiles where low income deciles from the mature economies were overrepresented. A striking fact is that among the percentiles in 1988 that turned out twenty years later to have been the most successful part of the global income distribution, 90% of people came from Asia. Among the 1988 percentiles that 20 years later turned out to have been the least successful part of the global income distribution, 86% came from mature economies.

The paper has offered another contribution. It created a new database, which (a) consists almost entirely of data derived from household survey micro data (b) keeps the household survey concept (income or consumption) constant over time for a given country, (c) uses only the most recent and most robust PPP exchange rates, and (d) allows a balanced panel analysis across the country-deciles. The latter aspect is particularly important since it has not been used before. In other words, while we knew that China’s growth rates were high and that the process was pro-rich (with higher deciles growing faster), we could not directly compare the growth of China’s top urban decile with (say) its 8th rural decile, and even less with the growth of deciles from the UK, Spain or Kenya. This is now all possible, and we have exploited here only a few aspects

opened up by the new database. Indeed we find that all Chinese, both rural and urban deciles, have improved their relative position in the global income distribution, at times dramatically, by jumping by more than 10 percentiles. The top urban Chinese decile thus went from being at the 68th percentile in the world to the 83rd. On the other hand, several African countries have experienced exactly the opposite evolution: the position of all or almost all of their deciles having gone down globally. China accounts for half of the top 20 most successful country-deciles in the period 1988-2008. Among the less successful country-deciles, the share of East European and more recently African deciles is preponderant.

References

- Abdelkrim, A. and J.-Y. Duclos: 2007, 'DASP: Distributive Analysis Stata Package'. PEP, World Bank, UNDP and Université Laval.
- Ackland, R., S. Dowrick, and B. Freyens: 2013, 'Measuring Global Poverty: Why PPP Methods Matter'. *Review of Economics and Statistics* 95(3), 813-824.
- Alessandria, G. and J. P. Kaboski: 2011, 'Pricing-to-Market and the Failure of Absolute PPP'. *American Economic Journal: Macroeconomics* 3(1), 91-127.
- Alvaredo, F.: 2010, 'The Rich in Argentina over the Twentieth Century, 1932-2004'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes: A Global Perspective*. Oxford University Press, pp. 253-298.
- Alvaredo, F., A. B. Atkinson, T. Piketty, and E. Saez: 2013, 'The World Top Incomes Database'. <http://topincomes.g-mond.parisschoolofeconomics.eu/>, Accessed 24/01/2013.
- Anand, S., P. Segal, and J. E. Stiglitz: 2010, 'Introduction'. In: S. Anand, P. Segal, and J. E. Stiglitz (eds.): *Debates on the Measurement of Global Poverty*. Oxford University Press.
- Anand, S. and P. Segal: 2008, 'What Do We Know about Global Income Inequality?'. *Journal of Economic Literature* 46(1), pp. 57-94.
- Armour, P., R. V. Burkhauser, and J. Larrimore: 2013, 'Deconstructing Income and Income Inequality Measures: A Crosswalk from Market Income to Comprehensive Income'. *American Economic Review: Papers and Proceedings* 103(3), 173-77.
- Atkinson, A. B.: 1970, 'On the measurement of inequality'. *Journal of Economic Theory* 2(3), 244-263.
- Atkinson, A. B.: 1975, *The Economics of Inequality*. Clarendon Press.
- Atkinson, A. B.: 2007, 'Measuring Top Incomes: Methodological Issues'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes over the Twentieth Century: A Contrast Between Continental European and English-Speaking Countries*. Oxford University Press.
- Atkinson, A. B. and A. Brandolini: 2010, 'On Analyzing the World Distribution of Income'. *The World Bank Economic Review* 24(1), 1-37.
- Atkinson, A. B. and T. Piketty: 2007, *Top incomes over the twentieth century: a contrast between continental European and English-speaking countries*. Oxford University Press.
- Atkinson, A. B. and T. Piketty: 2010, *Top incomes: a global perspective*. Oxford University Press.

Atkinson, A. B., L. Rainwater, and T. M. Smeeding: 1995, *Income distribution in OECD countries: evidence from the Luxembourg Income Study*. OECD social policy studies.

Banerjee, A. and T. Piketty: 2010, 'Top Indian Incomes, 1922-2000'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes: A Global Perspective*. Oxford University Press.

Bardasi, E., S. Jenkins, H. Sutherland, H. Levy, and F. Zantomio: 2012, 'British Household Panel Survey Derived Current and Annual Net Household Income Variables, Waves 1-18, 1991-2009'.

Barro, R. J. and X. Sala-i Martin: 1992, 'Convergence'. *Journal of Political Economy* 100(2), pp. 223-251.

Bhagwati, J.: 2004, *In defense of globalization*. Oxford University Press.

Bourguignon, F.: 1979, 'Decomposable Income Inequality Measures'. *Econometrica* 47(4), pp. 901-920.

Bourguignon, F.: 2011a, 'Non-anonymous growth incidence curves, income mobility and social welfare dominance'. *The Journal of Economic Inequality* 9(4), 605-627.

Bourguignon, F.: 2011b, 'A Turning Point in Global Inequality ... and Beyond'. Conference presentation, ABCDE conference.

Bourguignon, F.: 2012, *La mondialisation de l'inégalité*. Seuil et La République des Idées.

Bourguignon, F. and C. Morrisson: 2002, 'Inequality among World Citizens: 1820-1992'. *The American Economic Review* 92(4), pp. 727-744.

Burkhauser, R. V., S. Feng, S. P. Jenkins, and J. Larrimore: 2012, 'Recent Trends in Top Income Shares in the United States: Reconciling Estimates from March CPS and IRS Tax Return Data'. *Review of Economics and Statistics* 94(2), 371-388.

Carr-Hill, R.: 2013, 'Missing Millions and Measuring Development Progress'. *World Development* 46(0), 30-44.

CBO: 2012, 'The Distribution of Household Income and Federal Taxes, 2008 and 2009'. Congressional Budget Office Report 4441.

Chen, S. and M. Ravallion: 2004, 'How Have the World's Poorest Fared since the Early 1980s?'. *The World Bank Research Observer* 19(2), 141-169.

Chen, S. and M. Ravallion: 2010a, 'China Is Poorer Than We Thought, But No Less Successful in the Fight against Poverty'. In: S. Anand, P. Segal, and J. Stiglitz (eds.): *Debates on the Measurement of Global Poverty*. Oxford University Press.

Chen, S. and M. Ravallion: 2010b, 'The Developing World is Poorer than We Thought, But No Less Successful in the Fight Against Poverty'. *The Quarterly Journal of Economics* 125(4), 1577-1625.

Cowell, F. A.: 2009, *Measuring Inequality*. Oxford University Press.

Cowell, F. A. and E. Flachaire: 2007, 'Income distribution and inequality measurement: The problem of extreme values'. *Journal of Econometrics* 141(2), 1044-1072.

Deaton, A.: 2001, 'Counting the World's Poor: Problems and Possible Solutions'. *The World Bank Research Observer* 16(2), 125-147.

Deaton, A.: 2005, 'Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World)'. *The Review of Economics and Statistics* 87(1), pp. 1-19.

Deaton, A. and O. Dupriez: 2011, 'Purchasing Power Parity Exchange Rates for the Global Poor'. *American Economic Journal: Applied Economics* 3(2), 137-66.

Deaton, A. and A. Heston: 2010, 'Understanding PPPs and PPP-based National Accounts'. *American Economic Journal: Macroeconomics* 2(4), pp. 1-35.

Deininger, K. and L. Squire: 1996, 'A New Data Set Measuring Income Inequality'. *World Bank Economic Review* 10(3), 565-91.

Dynan, K. E., J. Skinner, and S. P. Zeldes: 2004, 'Do the Rich Save More?'. *Journal of Political Economy* 112(2), pp. 397-444.

Eltető, d. and P. Köves: 1964, 'On a Problem of Index Number Computation Relating to International Comparison.'. *Statisztikai Szemle* 42, 507-18.

Fisher, J. D., D. S. Johnson, and T. M. Smeeding: 2013, 'Measuring the Trends in Inequality of Individuals and Families: Income and Consumption'. *American Economic Review: Papers and Proceedings* 103(3), 184-88.

Frankel, D. M. and E. D. Gould: 2001, 'The Retail Price of Inequality'. *Journal of Urban Economics* 49(2), 219-239.

Gerschenkron, A.: 1947, 'The Soviet Indices of Industrial Production'. *The Review of Economics and Statistics* 29(4), 217-226.

Goldberg, P. K. and N. Pavcnik: 2007, 'Distributional Effects of Globalization in Developing Countries'. *Journal of Economic Literature* 45(1), pp. 39-82.

Grimm, M.: 2007, 'Removing the anonymity axiom in assessing pro-poor growth'. *The Journal of Economic Inequality* 5(2), 179-197.

Groves, R. M. and M. Couper: 1998, *Nonresponse in Household Interview Surveys*. Wiley.

- Haskel, J., R. Z. Lawrence, E. E. Leamer, and M. J. Slaughter: 2012, 'Globalization and U.S. Wages: Modifying Classic Theory to Explain Recent Facts'. *Journal of Economic Perspectives* 26(2), 119-40.
- Hlasny, V. and P. Verme: 2013, 'Top incomes and the measurement of inequality in Egypt'. World Bank Policy Research working paper 6557.
- Khamis, S. H.: 1972, 'A New System of Index Numbers for National and International Purposes'. *Journal of the Royal Statistical Society. Series A (General)* 135(1), pp. 96-121.
- Kopczuk, W., J. Slemrod, and S. Yitzhaki: 2005, 'The limitations of decentralized world redistribution: An optimal taxation approach'. *European Economic Review* 49(4), 1051-1079.
- Korinek, A., J. A. Mistiaen, and M. Ravallion: 2006, 'Survey nonresponse and the distribution of income'. *The Journal of Economic Inequality* 4(1), 33-55.
- Krueger, D. and F. Perri: 2006, 'Does Income Inequality Lead to Consumption Inequality? Evidence and Theory'. *The Review of Economic Studies* 73(1), pp. 163-193.
- Leigh, A. and P. van der Eng: 2009, 'Inequality in Indonesia: What can we learn from top incomes?'. *Journal of Public Economics* 93, 209-212.
- Li, H., L. Squire, and H.-f. Zou: 1998, 'Explaining International and Intertemporal Variations in Income Inequality'. *The Economic Journal* 108(446), 26-43.
- Milanovic, B.: 2002, 'True World Income Distribution, 1988 and 1993: First Calculation Based on Household Surveys Alone'. *The Economic Journal* 112(476), 51-92.
- Milanovic, B.: 2005, *Worlds Apart: Measuring International and Global Inequality*. Princeton University Press.
- Milanovic, B.: 2012, 'Global inequality recalculated and updated: the effect of new PPP estimates on global inequality and 2005 estimates'. *The Journal of Economic Inequality* 10, 1-18.
- Minoiu, C. and S. Reddy: 2012, 'Kernel density estimation on grouped data: the case of poverty assessment'. *The Journal of Economic Inequality* pp. 1-27.
- Mistiaen, J. A. and M. Ravallion: 2003, 'Survey compliance and the distribution of income'. Policy Research Working Paper, The World Bank p. 2956.
- Morival, E.: 2011, 'Top incomes and racial inequality in South Africa: Evidence from tax statistics and household surveys 1993-2008'. Master's thesis, Paris School of Economics.
- Nuxoll, D. A.: 1994, 'Differences in Relative Prices and International Differences in Growth Rates'. *The American Economic Review* 84(5), pp. 1423-1436.

Pinkovskiy, M. L.: 2013, 'World welfare is rising: Estimation using nonparametric bounds on welfare measures'. *Journal of Public Economics* 97(0), 176-195.

Pogge, T. W.: 2002, *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms*. Polity Press.

Rao, V.: 2000, 'Price Heterogeneity and "Real" Inequality: A Case Study of Prices and Poverty in Rural South India'. *Review of Income and Wealth* 46(2), 201-211.

Ravallion, M.: 2008, 'A Global Perspective on Poverty in India'. *Economic and Political Weekly* 43(43), pp. 31, 33-37.

Ravallion, M.: 2010, 'Understanding PPPs and PPP-based National Accounts: Comment'. *American Economic Journal: Macroeconomics* 2(4), pp. 46-52.

Ravallion, M. and S. Chen: 2003, 'Measuring pro-poor growth'. *Economics Letters* 78(1), 93-99.

Reddy, S. G. and T. Pogge: 2010, 'How Not to Count the Poor'. In: S. Anand, P. Segal, and J. E. Stiglitz (eds.): *Debates on the Measurement of Global Poverty*. Oxford University Press.

Sala-i Martin, X.: 2006, 'The World Distribution of Income: Falling Poverty and ... Convergence, Period'. *The Quarterly Journal of Economics* 121(2), 351-397.

Segal, P.: 2011, 'Resource Rents, Redistribution, and Halving Global Poverty: The Resource Dividend'. *World Development* 39, 475-489.

Shorrocks, A. and G. Wan: 2008, 'Ungrouping Income Distributions'. Working paper, UNUWIDER.

Singer, P.: 2002, *One World: The Ethics of Globalization*. Yale University Press.

Székely, M. and M. Hilgert: 1999, 'What's Behind the Inequality We Measure: An Investigation Using Latin American Data'. Research Department Working Paper Inter-American Development Bank.

Szulc, B.: 1964, 'Indices for Multiregional Comparisons'. *Przegląd Statystyczny* 3, 239-54.

van Kerm, P.: 2009, 'Income mobility profiles'. *Economics Letters* 102(2), 93-95.

Appendix

Appendix 1: Comparison with previous estimates of global inequality

Appendix Table 1: Comparison with previous estimates of global inequality					
	Benchmark years				
	1988	1993	1998	2003	2008
<u>Global inequality</u>					
<u>Gini index (%)</u>					
<i>Own estimates (Table 3)</i>	72.2	71.9	71.5	71.9	70.5
Milanovic (2012)	67.8	69.3	68.8	70.1	
Milanovic (2005)	61.9	65.2	64.2		
Milanovic (2002)	62.5	65.9			
Bhalla (2002) (Income)	67		65		
Bhalla (2002) (Consumption)	66		63		
Bourguignon & Morrisson (2002)		66			
Chotikapanich et al. (1997) (CVR)	65				
Dikhanov & Ward (2002)	69		68		
Dowrick & Akmal (2005) (GK)		64			
Dowrick & Akmal (2005) (Afriat)		71			
Sala-i-Martin (2006)	65	64	64		
Bourguignon (2012)	71		69		66
<u>GE(0) (Theil-L or mean-log-deviation) (%)</u>					
<i>Own estimates (Table 3)</i>	114.2	110.7	107.1	107.6	102.7
Milanovic (2002)	75.8	86.4			
Chotikapanich et al. (1997) (CVR)	80.6				
Dikhanov & Ward (2002)	102.1		97.1		
Sala-i-Martin (2006)	84.2	81.9	81.6		
<u>GE(1) (Theil-T) (%)</u>					
<i>Own estimates (Table 3)</i>	102.2	102.4	102.8	104.9	100.3
Milanovic (2012)	87.5	93.7	94.2	99.8	
Milanovic (2005)	71.5	81.8	79.2		
Bourguignon & Morrisson (2002)		85.5			
Dikhanov & Ward (2002)	89.1		90.7		
Dowrick & Akmal (2005) (GK)		79			
Dowrick & Akmal (2005) (Afriat)		101			
Sala-i-Martin (2006)	80.8	78.7	78.5		
<u>Decomposition by country</u>					
<u>GE(0) between-country contribution (%)</u>					
<i>Own estimates (Table 3)</i>	83.2	80.1	78.2	77.9	76.7
Milanovic (2002) (common sample)	75	74			
Sala-i-Martin (2006)	68	65	62		
<p><i>Notes: Milanovic (2012): Table 4, p. 14: Gini from row 5 (2005 PPP, sep. rural-urban prices for China, India & Indonesia); Theil from row 3 (2005 PPP, sep. rural-urban prices for China only); 2002 figures for 2003 benchmark. Milanovic (2002): Table 16, p. 72: Using full sample; Table 19, p. 78 (decomposition): Only for common sample. Milanovic (2005): Table 9.4, p. 108: Using full sample; Table 9.5, p. 112 (decomposition): Only for common sample. Bourguignon (2012): Figure 1, only approximate, because read-off from figure; 1988 refers to 1989, 1998 refers to 1997, 2008 refers to 2006. Otherwise: Anand and Segal (2008), Table 1: Survey estimates allocated to benchmark according to rules with micro data: 1988: Bhalla (2002), CVR (1997), and DW (2002) all refer to 1990; 1993: BM (2002) refer to 1992; 1998: Bhalla (2002) refers to 2000, and DW (2002) refers to 1999.</i></p>					

Appendix 2: Additional robustness checks for the top-heavy Pareto tail

In order to determine whether our results are sensible, we consider (1) the size of the national accounts adjustment, and (2) compare our Pareto coefficients to those in the World Top Incomes Database (WTID) (Alvaredo et al. 2013). The latter is important in order to find out whether our combination of allocating the entire gap to the top 10% and applying to such newly-calculated shares a Pareto adjustment produces Pareto coefficients similar to those observed in fiscal data in the WTID (which are thought to be reasonably good in capturing top income shares). Of course, we would not expect an exact correspondence because of differences (e.g. income definition or the unit of analysis) between the two data sources. Finally, we test for the robustness of the global Gini index to a restricted set of Pareto coefficients.

In order to summarise the magnitude of our national accounts adjustment, Appendix Table 2 presents the minimum and average values of the ratio (survey mean to mean used in the top-heavy Pareto adjustment) across the five benchmark years and across the different regions.⁸¹ Across the five benchmark years, the survey mean is on average between 75.1% and 81.4% of max(survey mean, national accounts consumption). The minimum value is 22.2% (Swaziland in 2008).

Appendix Table 2: Size of national accounts adjustment: Survey mean to max(survey mean, national accounts consumption) (%)											
	Benchmark years										1988-2008 change (pp)
	1988		1993		1998		2003		2008		
	mean	min	mean	min	mean	min	mean	min	mean	min	mean
World	81.4	41.3	80.5	25.4	80.8	22.5	78.8	36.9	75.1	22.2	-6.3
Mature economies	84.6	41.3	84.0	29.4	80.5	55.7	78.9	42.9	75.3	35.8	-9.3
China	100.0	100.0	90.6	90.6	95.3	95.3	92.2	92.2	100.0	100.0	0.0
India	77.5	77.5	71.7	71.7	60.0	60.0	56.5	56.5	53.0	53.0	-24.5
Other Asia	67.1	43.6	72.0	41.5	78.1	34.9	73.9	48.3	74.2	50.8	7.1
M. East & N. Africa	80.2	61.2	84.4	54.3	85.3	57.4	80.3	54.1	67.5	46.7	-12.7
Sub-Saharan Africa	91.1	52.6	80.0	25.6	84.5	22.5	87.1	53.2	81.4	22.2	-9.7
L. America & Carib.	80.1	48.5	75.7	25.4	81.3	43.4	73.0	45.2	73.3	45.6	-6.8
Russia, C. Asia, SE Europe	58.2	58.2	84.9	43.5	76.2	49.1	74.4	36.9	69.5	36.3	11.3

Notes: Observations are unweighted. Only for country-deciles which have National Accounts final household consumption.

We also match the surveys in our database with information from the WTID for the same survey year. For the country-years that we could match, the Pareto coefficient in the WTID ranges from 1.29 to 3.62 (over all the years). It is thought that the Pareto constant for income typically lies between 1.5 and 2.5 (Cowell, 2009). As explained in section 4, we have two Pareto imputations: the “proportional” and “top heavy” adjustments. The proportional adjustment tends to produce Pareto coefficients that are too high compared with the WTID, thus understating inequality (compare panel A of Appendix Table 3 below with panel C). This is not surprising given that we expect household surveys (from which the decile shares in this imputation are taken) to

⁸¹In other words, Appendix Table 2 shows the ratio of survey mean to max(survey mean, private national accounts consumption), which by definition is less than or equal to 1.

underreport top incomes. The Pareto coefficients obtained using the “top heavy” adjustment seem closer to the values in the WTID (compare panel B of Appendix Table 3 with panel C). However, the lowest values in this adjustment (just above 1) are implausibly low, thus probably overstating inequality. This is, as mentioned before, the result of the large national accounts discrepancies allocated to the top 10% entirely. Furthermore, only around 70% of the Pareto constants are within the range observed in the WTID, compared with 90% in the proportional adjustment.

Appendix Table 3: Comparison of Pareto constant with fiscal data					
	Benchmark years				
	1988	1993	1998	2003	2008
A. Baseline Pareto imputation (decile shares are unchanged)					
Mean	3.0	2.6	2.6	2.5	2.6
Median	2.7	2.5	2.4	2.4	2.6
Min	1.6	1.4	1.7	1.4	1.3
Max	12.8	4.6	4.1	4.2	4.3
Percentage within [1.29, 3.62] (range observed in WTID)	81.0	88.6	90.2	92.9	93.0
Number of surveys	63	105	112	126	114
B. Allocating National Accounts excess to top 10%					
Mean	2.0	1.8	1.8	1.7	1.7
Median	1.6	1.7	1.7	1.5	1.4
Min	1.1	1.1	1.1	1.1	1.1
Max	4.6	4.6	4.1	3.9	3.4
Percentage within [1.29, 3.62] (range observed in WTID)	74.6	77.1	77.7	77.8	68.4
Number of surveys	63	105	112	126	114
C. World Top Incomes Database					
Mean	2.5	2.4	2.1	2.0	2.0
Median	2.4	2.5	1.9	2.0	2.0
Min	1.5	1.4	1.3	1.5	1.6
Max	3.2	3.2	3.6	2.6	2.3
Number of surveys	14	18	21	18	12
<i>Notes: Observations are unweighted. Only for country-deciles which have National Accounts final household consumption. World Top Incomes Database: http://topincomes.g-mond.parisschoolofeconomics.eu/, accessed 24.01.2013.</i>					

These comparisons have to be taken with a note of caution. Of course, we would like Pareto constants (and thus top income shares) after the adjustment to be similar to the statistics obtained from WTID. But on the other hand, one has to acknowledge that the two databases refer to different income concepts (disposable vs. taxable income) and different units (individuals vs. taxable units). So, full correspondence between the two would be very unlikely. The problem however is that we have no yardstick to judge how close the two sources should ideally be.

In order to test the sensitivity of the global Gini coefficient to these implausibly low Pareto constants, we have decided to limit the bounds within which Pareto constants can lie. We consider two ranges: First, the range observed in the WTID, i.e. $a = [1.29, 3.62]$. Second, we chose tighter (essentially arbitrary) bounds such that $a = [1.5, 3]$. In both cases, the Pareto constants which are below the lower (above the upper) limit are changed to the lower (upper)

limit. Using these revised values of α , we calculate backwards the size of the national accounts gap and in the final step compute the revised shares and average fractile incomes.⁸² It might appear convoluted to adjust the Pareto constants and then compute the national accounts gap. The justification for proceeding this way (and we are fully aware that this is not a particularly strong reason) is that we have some guidance on what might be a sensible range of Pareto constants from the WTID.

As shown in Appendix Table 4, the effect of restricting the Pareto constants is that it reduces the global Gini by between 0.2 and 0.8 Gini points compared with “top heavy” adjustment from the main text (compare rows 1 and 2 in Appendix Table 4). For the tighter α limits of [1.5,3], the differences lie between 0.8 and 2.2 Gini points (compare rows 1 and 3). Thus, only the imposition of tighter limits on the admissible Pareto constants may have a sizeable impact on the global Gini. Nevertheless, compared with the baseline results using survey means, the global Gini index is still substantially greater.

Appendix Table 4. Further robustness check on the global Gini index:
imposing limits on the Pareto constants

	Benchmark years					1988-2008	1993-2008
	1988	1993	1998	2003	2008	change (pp)	change (pp)
(1) Private consumption with top-heavy Pareto imputation 1/	76.3	76.1	77.2	78.1	75.9	-0.5	-0.2
(2) Private consumption with top-heavy Pareto imputation and broader α limits	76.1	75.8	76.7	77.5	75.1	-1.0	-0.7
(3) Private consumption with top-heavy Pareto imputation and tighter α limits	75.6	75.2	76.2	75.9	74.3	-1.3	-1.0
Number of surveys	63	105	112	126	114		

1/ From Table 4.

Notes: Observations are weighted using population. All calculations are done across the sample of 520 surveys.

⁸² The national accounts adjustment (ϵ) is given by

$$\epsilon = \frac{1}{10} \left[\frac{\bar{Y}_9}{2^{\frac{\alpha-1}{\alpha}} - 1} - \bar{Y}_{10} \right]$$

where \bar{Y}_9 and \bar{Y}_{10} are the average incomes of the 9th and 10th decile observed in the household survey.

Appendix 3: The winners and losers between 1993 and 2008

Appendix Table 5: Winners and losers in terms of average annualised growth (1993-2008)			
20 biggest gainers		20 biggest losers	
(best at the top)		(worst at the bottom)	
decile	growth	decile	growth
Swaziland-1	19.1%	Japan-3	-2.1%
Panama-1	17.7%	Burundi-7	-2.1%
Swaziland-2	16.3%	Slovak Rep.-1	-2.1%
Kenya-1	16.2%	Bulgaria-1	-2.2%
Lithuania-1	15.9%	Israel-1	-2.4%
Romania-10	15.4%	Centr. Afr. Rep-3	-2.5%
Swaziland-3	14.2%	Burundi-6	-2.5%
Azerbaijan-2	13.9%	Kyrgyz Rep.-10	-2.7%
Azerbaijan-1	13.8%	Japan-2	-2.9%
Azerbaijan-3	13.8%	Burundi-5	-3.1%
Azerbaijan-4	13.7%	Centr. Afr. Rep-2	-3.3%
Romania-9	13.5%	Burundi-4	-3.9%
Azerbaijan-5	13.4%	Japan-1	-4.5%
Azerbaijan-6	13.2%	Burundi-3	-5.1%
Romania-8	13.1%	Bolivia-1	-5.3%
Azerbaijan-7	13.0%	Kenya-10	-5.8%
Swaziland-4	13.0%	Honduras-1	-6.3%
Romania-7	12.8%	Centr. Afr. Rep-1	-6.5%
Lithuania-2	12.7%	Burundi-2	-7.3%
Romania-6	12.7%	Burundi-1	-13.1%

Notes: Only for countries observed in 1993 and 2008 and which have at least one 5-year growth interval. Deciles are numbered 1 to 10, with 1 being the bottom decile.