Working Paper Series

Inequality from a global perspective: An alternative approach

Peter Edward
Andy Sumner

ECINEQ WP 2013 - 302

ECINEQ 2013 - 302
August 2013

www.ecineq.org

Inequality from a global perspective: An alternative approach*

Peter Edward
Newcastle University Business School

Andy Sumner
King's College London

Abstract

The interplay of between- and within-country inequality, the relative contribution of each to overall global inequality, and the implications this has for who benefits from recent global growth (and by how much), has become a significant avenue for economic research. Drawing conclusions from aggregate inequality indices such as the Gini and Theil reduces the highly complex nature of global inequality to a single coefficient and makes it impossible to take a nuanced view of how global growth interacts with changing national and international inequality. In light of this we propose and justify an alternative approach based on four consumption 'layers' identified by reference to the global consumption distribution. We consider how each 'layer' of global society has fared since the end of the Cold War.

Keywords: Poverty, Inequality, Economic Development.

JEL Classification: D63, I32.

*We would like to thank the following for comments on earlier drafts: Nora Lustig, Stephanie Seguino, Borge Wietzke, Martin Evans, Simon Maxwell, Owen Barder and Richard Manning.

Glossary

EAP	East Asia and the Pacific
GDP	Gross Domestic Product
GNI	Gross National Income
GrIP	The 'Gr'owth, 'I'nequality and 'P'overty Model
HFC	Household Final Consumption
HIC	High-income country
LAC	Latin America and the Carribbean
LIC	Low-income country
LMIC	Lower middle-income country
MLD	Mean-log deviation
NA	National Account means
pa	per annum
pc	per capita
PPP	purchasing power parity
S	Survey means
SAR	South Asia Region
SSA	sub-Saharan Africa
UMIC	Upper-middle income countries
UNU	United Nations University
WDI	World Development Indicators
WIID	World Income Inequality Database

1. INTRODUCTION

The interplay of between- and within-country inequality, the relative contribution of each to overall global inequality, and the implications this has for who benefits from recent global growth (and by how much), has become a significant avenue for economic research. In this paper we provide new estimates of the evolution of inequality between and within countries and explore who have been the winners and losers from global growth in recent decades. Estimates are derived using a custom-built model of economic growth, inequality and poverty.

The paper is structured as follows: Section 2 reviews relevant recent empirical studies. Section 3 outlines the model. Section 4 asks what has happened to between- and within-country inequality over the last 20 years using the Gini and Theil and density curves. Section 5 then presents an alternative, 'layers' approach to inequality. Section 6 concludes.

2. GLOBAL INEQUALITY AND GROWTH

A search of peer-reviewed studies published in economics and development journals since 2000 identified more than 120 papers relating to either the empirical study of long-run trends in income/consumption inequality or the distribution of the benefits of economic growth.[1] Different studies have used different concepts, methods and datasets and are typically non-comparable in any meaningful sense as they explore

[1] These include the oft-cited Bourguignon and Morrison (2002), Dowrick and Akmal (2005), Milanovic (2012b), Sala-i-Martin (2006) and Wade (2004). The search criteria was as follows: (i) empirical in nature; (ii) long-run in nature – meaning they span periods of at least five years; (iii) based on cross-country analysis (rather than studies of one or a small number of countries); and (iv) published since 2000 (a suitable cut-off because of the improvements in inequality data that emerged at the end of the 1990s.

inequality differently.[2] Evidently there is a high level of interest and concern in the ways that growth, inequality and poverty interact.

Several noteworthy recent papers deal with 'class' (defined in various ways), 'geography' (meaning location) and changes in inequality since 1990 in order to sustain arguments about changes in inequality and the distribution of the benefits of growth during that period. These papers include the following which in particular have been conceptually and methodologically innovative: Milanovic (2011, 2012a), Palma (2011) and Ravallion and Chen (2012).

These studies have made two arguments in particular: first, that 'global inequality' (defined in different ways) is falling because international (between countries) inequality is falling. Second, that within-country inequality is rising in fast growing Asia, albeit from relatively low levels, and is falling in Latin America, albeit from very high levels. Elsewhere, trends in within-country inequality in sub-Saharan Africa (SSA) are difficult to discern regionally with clarity.[3] Significant issues are therefore the interplay of between- and within-country inequality and the dominant role of China in the dataset. One question is thus if China is excluded, then from 1990 to 2010 what has happened to between- and within-country inequality across the rest of the world?

The papers noted above approach between- and within-country inequality in different ways. For example, Palma (2011) focuses predominantly on similarities and differences in within-country inequality across the world. Using a dataset that includes observations for 135 countries (ibid. p.89) with information on Gini and

[2] See reviews of Sudhir and Segal (2008) and Goda (2013) for further discussion of a range of concepts and methods in papers.

[3] Studies on Latin America, argue that the changes in inequality are partly policy induced (for example by cash transfer programmes) and partly the result of a fortuitous interplay of changes in relative wages for skilled and unskilled workers due to commodity booms creating higher demand for relatively unskilled labour (see Lustig et al., 2011; 2013).

income shares from World Development Indicators (World Bank, 2012), Palma reaches the following conclusions: first, he shows that about 80 percent of the world's population now live in regions whose median country has a Gini close to 40, implying a reduction in regional differences in within-country inequality. Second, that there are two opposite forces at work. One is 'centripetal' and leads to a growing uniformity in the income-share appropriated by the 'middle' 50 percent (deciles 5 to 9). The other is 'centrifugal' and leads to an increased diversity in the shares of the top 10 percent (decile 10) and bottom 40 percent (deciles 1 to 4). The share of the national consumption/income that accrues to, or is appropriated by, the five 'middle' deciles (deciles 5 to 9) is remarkably constant across countries, with a median value of 52 percent and most values within the range of 50-55 percent (Palma, 2011, p. 101, 102). This, he argues, indicates that, regardless of differences in national per capita income or in political-institutional settlements:

> half of the world's population (the middle and upper-middle classes) have acquired strong 'property rights' over half of their respective national incomes (ibid, p. 103–104).

The remaining half of the national income is variously distributed between the poorest 40 percent (deciles 1 to 4) and the richest 10 percent (decile 10) in each country. There is, however, a wide diversity, between countries, in the share appropriated by the richest 10 percent. This leads Palma to propose that the problem of national inequality is one of 'Homogenous Middles, Heterogeneous Tails' in which the key

issue is the capture of GNI of the richest and the poorest.[4]

While Palma focuses on within-country inequality, Milanovic (2011, 2012a) in contrast, uses data (2011, p. 7) derived from the World Income Distribution database[5] to combine within-country inequality distributions with between-country inequality measures (mean per capita consumption in real PPP $, based to 2005) to construct a model of inequality between all individuals in the world.[6]

Milanovic's (2012a) central argument derived from the 'Inequality 3' concept (see footnote 6) is that:

> perhaps for the first time since the Industrial Revolution, there may be a decline in global inequality. … [i]t is indeed among the very top of the global income distribution and among the "emerging global middle class", which includes more than a third of world population, that we find most significant increases in per capita income. (2012a, pp. 7–8, 12).

[4] One interesting observation from this is that a country's Gini coefficient is therefore predominantly dependent on the income share of the richest decile. Palma notes that the Gini coefficient can be (roughly) estimated as 1.5 times the share of the top 10 percent (in percentage points) minus 15, which implies (if one assumes that the share of the lowest 40 percent is negligible and that inequality among Palma's 'middle' class is also negligible) that the maximum likely Gini would be in the region of 60 percent (2011, p. 103). Notably, Palma identifies two middle-income regions – Southern Africa and Latin America – where countries typically have Gini indices between 55 and 60, in other words very close to the maximum likely Gini value. In practice a few Gini indices slightly higher than 60 can be found (e.g. Seychelles 66, Comoros 64, Namibia 64, South Africa 63, Micronesia 61, Botswana 61). Of course Gini indices higher than 60 are possible when we look at global or regional inequality because this adds in the effect of between-country inequality.

[5] The dataset is available on http://econ.worldbank.org/projects/inequality

[6] Milanovic calls this 'Inequality 3'. By contrast, Inequality 1 is a measure of between-country inequality in which all countries carry equal weight. Inequality 2 is also a between-country inequality measure, but it is population weighted. As Milanovic notes, it is the vastly improved coverage of global data (both national distributions surveys and international PPP comparator data) since the late 1980s that has enabled analysts to move to 'truly global' (i.e. Inequality 3) modelling, and in doing so to start to consider in some detail how and where growth and inequality interactions lead to winners and losers in the global economy. By the early 2000s coverage of the data meant that some initial models of true global inequality started to emerge. In essence, all of these can be understood as Inequality 3 types of models, albeit often with not insignificant differences in assumptions and results, notably from how they modelled disaggregated national distributions and how they assigned relative weightings when aggregating country-data into a global picture. Examples of these include Bhalla (2002), Dikhanov and Ward (2002), Edward (2006), Milanovic (2002, 2005), and Sala-i-Martin (2002).

Milanovic (2012a, p. 18) also reprises Milanovic (2011) and decomposes global inequality between 'class' ('differences in incomes within nations') and 'location' ('differences between mean incomes of all the countries in the world') and notes that:

> Around 1870, class explained more than 2/3 of global inequality. And now? The proportions have exactly flipped: more than 2/3 of total inequality is due to location (2012a, p. 19).

In short, Milanovic argues that, compared to the late nineteenth century, there is much more inequality between countries today but less inequality within countries.

Milanovic's use of the term 'class' is questionable. In general he, like Palma, sees class as a national issue (as measured by within-country inequality) but he is also happy to refer to an 'emerging global middle class' that is defined by reference to location in the global income ('Inequality 3') distribution.[7] Milanovic highlights the benefits of disaggregating the ('Inequality 3') global income distribution to look at how different 'segments' of the global population (delineated either in absolute terms or relative to the global distribution) have fared during the last two decades.[8]

Finally, Ravallion and Chen (2012) discuss changes in inequality utilising the 'mean-log deviation' (MLD) and the PovcalNet database (of more than 850

[7] From Milanovic's Figure 4 (2012a, p. 13), this 'class' appears to be the global population between the 25th and 60th percentiles.

[8] While Milanovic does not define explicitly the 'global middle class', Ravallion (2010) does do so by per capita expenditures (as does Kharas, 2010). He also notes that there is little agreement over what these limits should be: 'Milanovic and Yitzhaki (2002) defined it as the set of people living between the mean incomes of Brazil and Italy. Instead Banerjee and Duflo (2008) defined the middle class as those living between $2 and $10 a day at 1993 PPP. Bhalla (2007) ... proposed a lower cut-off point of ... about $10 per day in 2005 [PPP] ... he set an upper bound at 10 times his lower one." (Ravallion, 2010, p. 446).

household surveys from 127 developing countries) to explore changes in within- and between-country inequality from 1980 to 2008. The MLD is 'the difference between the log of the group's mean consumption and the mean of the logs of all the consumptions within that group' (ibid., p. 2). Unlike the more commonly used Gini coefficient, the MLD can be decomposed by population sub-groups so that the population-weighted MLD provides an estimate of the contribution of within-county inequality to overall inequality. Ravallion and Chen (2012, p. 2) note,

> We see that there has been a trend decrease in total inequality, though with ups and downs, and an increase over 2005-08. However, that pattern has largely been due to inequality between countries…. Over the period as a whole, we see that [the within-country component] has risen... The within-country component accounted for less than one-third of inequality in the developing world as a whole in 1981, but almost half in 2008.

In sum, there are a set of disagreements in these recent long-run attempts to consider global and national inequality and the relative importance of each. Whereas Milanovic suggests that from 1870 to 2000 within-country inequality ('class') became less important and between-country inequality ('geography') more important, Ravallion and Chen (2012) find the opposite trend over the period 1981–2008, namely that since the early 1980s within-country inequality has become an increasingly important element in overall global inequality.[9] It seems possible therefore that the long-run trend of the twentieth century towards lower inequality within countries but higher inequality between countries could be starting to reverse,

[9] Ravallion and Chen (2012, p. 4) also note significant variations between regions.

with the extremes between rich and poor people increasing within countries while the income gaps between rich countries and poor countries start to reduce. However, what if these differences can be largely attributed to a few large emerging economies, most notably China, which disguise that large differences between the remaining countries remain rather intractable?

3. A MODEL OF GROWTH, INEQUALITY AND POVERTY

In order to assess the debate above, this paper makes use of a custom-built model of growth, inequality and poverty. Henceforth referred to as the GrIP ('Gr'owth, 'I'nequality and 'P'overty) model (version 1.0). The core approach in the GrIP model is to take for each country the distribution (quintile and decile) data and, by combining this with data on national population and on the mean consumption per capita in internationally comparable PPP $, develop for each country an estimate of how many people live at any specific consumption ($-a-day) level. This is very similar to the approach in Povcal except that GrIP uses a linear rather than kernel distribution function and draws on a wider range of sources to extend the analysis to include developed economies. Having identified for each country the number of people living at a given consumption level, GrIP then aggregates these to build a truly global income distribution (of course, a wide variety of other aggregations are also readily produced; for example, by region or income category as shown in the various results presented here). These aggregations can then be interrogated to investigate issues such as poverty levels, trends in inequality and who are the winners and losers from global growth.

Survey distributions (quintile and upper and lower decile data) are taken (in this order of preference) from PovcalNet,[10] World Development Indicators or the UNU WIID V2.0c (May 2008) database.[11] National consumption means can come from survey data or from National Account (NA) measures. In the discussion for this paper we make use of survey means to remain comparable to the reference literature.[12] All analysis and results are in 2005 PPP $.

There are a number of methodological issues to consider. First, even though these datasets have greatly improved their global coverage in recent years, there are still some significant gaps in the data. Surveys do not take place annually so in the GrIP model distributions for intermediate years, between-surveys are calculated by interpolation, while in years subsequent to the most recent survey the distribution is assumed to remain unchanged from that survey. Where a country has no surveys, or the gaps between surveys are too great to allow reliable interpolation, the GrIP model can 'fill' a country's missing distributions with the (not population-weighted) average distribution from all other countries in the same region and income group. This means that the analysis can either be 'filled' to more closely replicate global population and consumption totals, or 'not filled' to include only the smaller set of countries for which national distribution data is available.[13]

The GrIP model also, and unlike other models, disaggregates the national populations into globally standard '$ per capita' brackets, thereby avoiding

[10] In this paper we use Povcal version of February, 2012.

[11] See www.wider.unu.edu/research/Database/en_GB/database. Where WIID V2.0c is used consumption distributions are used in preference to income distributions. In accordance with established practice, no attempts are made to modify income distributions to 'convert' them to consumption distributions on the basis that such conversions are too speculative to be justified.

[12] Edward and Sumner (2013) compare NA and survey means with reference to global and regional poverty estimates.

[13] We note also that the distribution data can be derived at either the individual level or the household level. This is an outcome of the original survey design and so it is difficult to adjust for in subsequent analysis. As is the case for most other studies we do not attempt to adjust for this difference but note that household surveys will inevitably understate national inequality to some extent as they do not include intra-household inequality.

introducing the distortions of approaches, such as Bhalla's simple accounting procedure (Bhalla, 2002; Hillebrand, 2008), where by disaggregating only to percentiles some large step-change distortions are introduced in the later global aggregation at points where percentiles from the very largest countries (such as India and China where each percentile currently includes well over 10 million people) are added back into the global distribution.

In this paper we use Household Final Consumption (HFC) per capita means (in 2005 PPP $). Where GDP data exists but HFC data does not then the missing HFC data is estimated from the GDP data. Table 1 illustrates how by first estimating missing HFC data from GDP data (for countries that otherwise have valid survey distributions) and then using filling to estimate distributions for countries without valid surveys, the GrIP model incrementally builds a global model of inequality from the available source data. It can be clearly seen that the number of countries underpinning the model, and hence also the reliability of any outputs from the model, reduces rapidly once we go back into the 1980s. For this reason the results given here mainly focus on the period from 1990 to 2010. Where we do take analysis back into the 1980s those results should be treated with circumspection.

We considered a set of sensitivity tests as follow: survey versus NA means, GDP vs HFC, and filled or not filled analysis. We chose HFC not GDP because it was closer to actual consumption as measured in surveys. We have shown some of the NA vs survey differences in the text but in general only provide survey means because it makes our work comparable to the reference literature. And we provide rationale in the text for use of filled or not filled at different points.

Table 1: Coverage (cov.) of analysis and effects of estimating HFC and filling distributions

4. GLOBAL INEQUALITY AND DECOMPOSITIONS

4a. Inequality by Gini and Theil measures

Two commonly used measures for global inequality are the Gini index and the Theil index. Of these, the Gini is the more widely used, largely because of its close and relatively intuitive association with the Lorenz curve. However the Gini is not decomposable so that it does not unambiguously differentiate the separate contributions of within- and between-country inequality (it includes a significant 'overlap' or 'interaction' term between the within- and between-country contributions). The Theil index is, however, fully decomposable, but as a measure of entropy it is rather less intuitive but importantly, is generally more sensitive to changes at the extreme ends of the Lorenz curve, in comparison to the Gini which is more sensitive to changes in the middle of the distribution (see for full discussion, Cowell, 2000).

We provide estimates here using both indices for the period from 1980 to 2010 (all based on 'not filled' analysis so that estimates are not distorted by any estimations for 'missing' countries) and include for comparison estimates from Milanovic (2012b, p. 14). Many other earlier estimates of Gini and Theil exist but they are not directly comparable to our analysis because, as Milanovic (2012b) shows, the 2008 rebasing of international PPP rates systematically increased global inequality indices.

Looking across the figures below it is apparent that global inequality has been falling since 2000. It is less clear whether inequality was static (see Gini figures) or rising (see Theil figures) in the 1980s and 90s (and the change in the Theil and lack of change in the Gini might point towards changes taking place at the extremes rather than the middle of the distribution). What is clear, however, is that since the mid-1980s within-country inequality has risen slowly but steadily. However, since the early 2000s between-country inequality has started to fall more quickly (although that does not guarantee that these falls will continue into the future) so that overall global inequality has also started to fall. Taken as a whole the results indicate that for the first ten to fifteen years since the late 1980s, global inequality was relatively static with a slow but steady rise in within-country inequality being partly offset by a gradual decline in between-country inequality. Since then, and particularly since 2005, while within-country inequality has continued to rise, albeit perhaps rather more slowly, between-country inequality has fallen quite rapidly, and with it global inequality has started to fall too. The interaction of these effects means that whereas in 1988 within-country inequality accounted for around 20 percent to 25 percent of global inequality, by 2010 it had risen to 30 percent of global inequality.

The figures below are consistent with Ravallion and Chen (2012) and appear to be the reverse of the longer term trend Milanovic identifies since 1870 (2012a, p. 18). It is possible that our model is detecting the first signs that the world is trending back towards the situation in the past where countries are more equal relative to each other but more unequal within themselves.[14]

However, the picture looks rather different when China is excluded (comparing Figures 1, 3 and 5 which include China with Figures 2, 4 and 6 which

[14] Milanovic estimates that in 1870 the global Theil Index was about 65, with two-thirds of global inequality being due to within-country inequality.

exclude China). In the rest of the world within-country inequality has overall been remarkably constant – as some countries have become less equal others have become more so. But between-country inequality rose steadily in the 1980s and 1990s. Since 2000 between-country inequality has been fairly static (when China is excluded). On the basis of what we see here, it would seem that recent falls in global inequality are predominantly attributable to rising prosperity in China. Elsewhere, a trend beginning in 1980 of increasing inequality between rich countries and poor countries may have stalled since 2000 but it is not apparent that it has gone into decline even after the 2008 financial crisis. This would suggest that very modest signs of falling overall global inequality are due to the rapid progress of China which may be masking underlying trends that are quite different.

Figure 1: Global Gini coefficient, survey means (not filled)

Figure 2: Global Gini coefficient excluding China, survey means (not filled)

Figure 3: Global Theil Index, survey means (not filled)

Figure 4: Global Theil Index excluding China, Theil Index, survey means (not filled)

Figure 5: Within-country Theil component as percentage of global Theil Index, survey means (not filled)

Figure 6: World excl. China, Within-country Theil component as percentage of global Theil Index, survey means (not filled)

4b. Inequality by density curves

Global distribution curves show that in the mid-1980s, with caveats noted earlier, we lived in what was predominantly a ‘twin-peak’ world (see Figure 7). In other words

there was a fairly well-defined global distinction between a substantial 'poor' peak and a smaller 'rich' peak (to which accrued most of global GDP and consumption). This gives the 1985 density curve a distinct 'dumb-bell' shape, indicating the division of world population into two fairly well-defined and distinct segments – the old North-South or West-Rest division. Since then the dumb-bell has become much less distinct (see Figure 8) (best seen as the loss of concavity between the poor and rich peaks in the density curves, presented below, so that the dumb-bell looks, in 2010, more like a rotated rectangle). [15]

Figure 7: The 'twin-peak' global distribution of the mid-1980s

Figure 8: The 2010 global distribution and loss of the 'twin peaks' or 'dumb-bell'

It would be tempting to say that this indicates a dismantling of the twin-peak world and the formation of a single more inclusive global cluster (the formation of a single population peak). However, as the figures excluding China show, closer inspection reveals that that may not be what is happening since once China is removed, the population curves remain concave at consumption over $2 a day. Another way of interpreting this is that in the twin-peak 'dumb-bell' world of the 1980s the world had a missing 'global middle'. What has happened in the last two decades is not that this old global distribution has been radically disrupted but that the missing middle has started to be filled in. That the old 'dumb-bell' structure largely persists is indicated in the 2010 curves which still clearly show a distinct peak for the global poor (best seen

[15] The old twin-peaks world was identified by Quah (1996). The likelihood of the trend away from a two-peak world was previously detected in Edward (2006, p. 1677). As noted earlier, the recent 'consumption' density curves also show the emergence of a new peak at the highest consumption levels; that is, approaching $100,000 PPP pc pa. This was present back in 1985 but seems to be becoming more distinct recently – perhaps indicative of the emergence of a 'super-rich' segment. However, this phenomenon needs to be treated circumspectly for reasons discussed in the text.

on the population curve, albeit now at an income near the $2 a day line whereas in 1990 it was close to $1.25) and another distinct peak for those living at around or above $50/day (best seen on the income curve). A peak for a reasonably homogeneous cluster would be expected to have more of a normally distributed bell-curve shape, whereas the 2010 population curve if viewed as representing a single cluster displays a strongly skewed shape (which would be all the more evident if the horizontal $ per capita/annum axis was not plotted on a log scale). In other words, rather than seeing the formation of a more globally integrated socioeconomic structure (one where distinct differences in consumption patterns become superseded by a more continuously graduated range of consumption levels dispersed across all the world), we may merely be witnessing the emergence of a more complex structure where (temporarily at least) the missing middle between the twin peaks is being filled by an emerging global 'middle'. The slight peak emerging around the 2010 $10 a day level might also indicate that this middle is developing into a distinct consumption 'layer', although this could just be a transition stage as some people in emerging economies move across.

That this emerging middle is unlikely to represent an end to the strong separation of the rich and poor clusters of the old twin-peak world can be illustrated by looking at the various country income categories separately. Figures 9 to 15 show density curves for current (2010) categories of high-income countries (HICs), upper middle-income countries (UMICs), and low-income plus lower middle-income countries (LICs plus LMICs). These diagrams clearly reveal that the aggregated distributions across the *current* LICs and LMICs still retain the reasonably well-defined and balanced forms that they had back in 1990. In other words, the growth of the last two decades has not radically altered the location of the peoples of these countries as the

global poor. Collectively, their population and consumption distributions display the well-defined forms that might indicate a reasonably homogeneous global consumption grouping. Indeed, if anything, it looks as if the distributions in these countries have become more normal, which would only strengthen their claim to represent a reasonably homogeneous global consumption clustering (see later discussion). Noteworthy also is that, despite overall economic growth and rising mean incomes, total numbers for extreme poverty ($1.25 a day) across the LICs and LMICs have not changed much since 1990.

Figure 9: Density curves for current HICs

Figure 10: Density curves for current LICs and LMICs

Figure 11: Density curves for current UMICs

Figure 12: Density curves for current LICs and LMICs excluding India

Figure 13: Density curves for current UMICs excluding China

Figure 14: Density curve for India

Figure 15: Density curve for China

If we look at the high-income countries there is little evidence of radical changes in those living in the range of $10–$50/day of the distribution structure here (that is, in the region that would be affected if we were witnessing global convergence away from the twin-peak structure). The high-income countries still largely retain the presence of a relatively homogeneous ‘global secure’ component with the main changes in these countries being some evidence of the evolution of greater differentiation among those living above the $50/day level in these countries.

It is really only in the picture for the upper middle-income (UMIC) countries, which include China, that we can detect radical changes in the shapes of the curves as these countries move in to fill the missing global middle. Here we do see not only some rapid changes in income levels (the shape of the upper 'population' curve) and in the concentration of buying power (the shape of the lower 'income' curve) but also evidence for the emergence of a number of mobile and overlapping density peaks.

In other words, most of the 'structural' change in the distribution of global consumption is confined to the UMICs. If we omit the UMICs then the long-standing twin-peak 'dumb-bell' of a bipolar world is seen to persist (Figure 16). Leaving aside the putative issue of the very-rich peak, we can see that the HICs retain a strong homogeneity in their consumption clustering in the sense that there is a fairly well-defined peak to which most of their society belongs.

Figure 16: Density curves for all countries excluding current UMICs

Similarly in the aggregated LMICs and LICs there is still a reasonably well-defined peak. There is some evidence that wealthier elements may have been pulling away in the 1990s (the incipient formation of a second population peak in 1995) but the 2010 curve seems to be returning to a more homogeneous form (closer to a single bell curve). The 1985 global curves clearly show that at that time we lived in a very dichotomous world with a low-income peak and a high-income peak separated by a relatively 'thin' middle. Since then at the global level this division seems to have become less distinct but this apparent loss of global separation is really only confined to the UMICs. In the rest of the world, the long-standing global structure of inequality persists with a relatively rich grouping living in the high-income countries (and

clustered around a median consumption of about $25/day per capita in 2005 PPP international $) and a relatively poor grouping living in the LMIC/LIC countries (and largely clustered around a median consumption of about $2/day per capita in 2005 PPP international $ or more precisely $1.70 per capita).

Certainly within China we can see the evolution from a predominantly poor society in 1990, with most of the population living around a peak centred just below the $1.25 extreme poverty line, to a much more heterogeneous society in 2010, in which there are three distinct peaks with a low consumption cluster centred on $3 a day, a smaller middle cluster centred on $6.50 a day and a richer cluster centred close to $18 a day. The data suggests heterogeneity is mostly increasing in the UMICs while LICs and LMICs with or without India have if anything become more homogeneous (a clearer single peak).

5. AN ALTERNATIVE APPROACH TO GLOBAL INEQUALITY

5a. Inequality beyond aggregate indices

Evidently, exclusion of a single (albeit large and rapidly growing) country can dramatically modify the picture on what is happening to global inequality. Additionally, inequality coefficients demonstrate some significant national differences in trends over time that quickly become lost sight of when ‘rolled-up’ into a single global inequality measure (see, for example, Figure 17).

Figure 17. Gini coefficient, 1980–2010 in selected populous countries

One conclusion to draw is that by reducing the highly complex nature of global inequality to a single coefficient it becomes difficult to then take a nuanced view of how global growth interacts with changing national and international inequality. An alternative approach would be to identify global absolute consumption 'layers' and explore the changing fortunes of each 'layer'. We introduce density curves which illustrate, across the full spread of global per capita consumption levels, both the distribution of global population at each consumption level and their share of global consumption. We then present a range of fractile curves which illustrate, across global, regional and national populations, which per capita consumption levels benefited most from global growth (in terms of their percentage growth in consumption) since 1990. Together these build a much more nuanced picture.

5b. The four 'layers' of global society

In this section we identify, with reference to peaks in the global consumption distribution, four global 'layers' as follows by per capita consumption: the 'global absolute poor'; the 'global prosperous' and those in-between, specifically, the 'global insecure' and the 'global secure'.

This approach, of segmentation by absolute consumption level, has been used to estimate global dollar-a-day poverty levels. But those analyses focus only on the poorer countries and only on the lowest income levels (numbers below a global absolute poverty line). More recently, a body of empirical studies related to developing countries has emerged in response to the growing data on middle-income groups. Typically referring to these groups as 'the middle classes', more often than not, the segmentation is defined by reference to daily expenditure per capita. Many of

these recent studies are based on absolute definitions of expenditure per capita/day (PPP), ranging from $2/day to $100/day (see review of Sumner, 2012). For example, Ravallion (2010) takes the segment between a lower-bound absolute threshold of $2 a day and an upper-bound threshold of $13 a day (the US poverty line) (all in 2005 PPP $). However, this work is based on Povcal and so only covers developing world countries. Kharas (2010) too has defined the middle class against absolute thresholds with a lower bound of $10 and an upper bound of $100.

Further, there have been some attempts in sociology to identify global classes. These have largely been limited in recent years to considering whether there is now a very small but distinct class of 'transnationally-oriented elites grounded in globalized circuits of accumulation' (Robinson, 2012). However, even in that model this transnational class still competes with nationally oriented classes, which include both local elites and other popular and working classes with strong national identities. And even these theories, which see a very specific and limited scope for the notion of global classes, are strongly disputed (Carroll, 2012).

Class is, of course, often discussed in terms of types of assets and productive processes, labour markets, and occupational resources (see discussion in Torche & López-Calva, forthcoming). Some contemporary sociological analysis of class also places a particular emphasis on economic security (see, for discussion, Goldthorpe & McKnight, 2006; Standing, 2011). However, the conflation of per capita expenditure and class, while not without some merit, is difficult to sustain because 'class' is a social and political identity not necessarily linked to estimates of expenditures per capita.[16] In short, any segmentation by absolute consumption levels groups people

[16] Note also that Palma (2011) and Milanovic (2012b) understand class as a national issue implying that class is more relevant to national differences than to global ones and therefore that class segmentation should be relative not absolute. This is also why their categorisations do not translate into a system of global absolute segmentation.

globally by their consumption levels rather than by any deeper 'class'-derived alignment of sociopolitical orientation. The precedent for segmentation by consumption level lies less in class theory than in preference similarity theory – the idea that people with similar purchasing power levels tend, wherever they are in the world, to have broadly similar consumption preferences (Linder, 1961).

We propose four consumption 'layers' (not 'classes') of global society: the 'global absolute poor' ($0-$2 per capita per day); the 'global insecure' ($2-$10); the 'global secure' ($10-$50); and the 'global prosperous' ($50+). We therefore identify three thresholds between these layers, namely $2, $10 and $50 divisions (all in terms of survey means). Identifying such 'benchmark' thresholds is inevitably a rather rough and ready exercise but since we are applying them to a global consumption distribution it makes sense to derive them in relation to the patterns of that distribution. By looking at consumption distribution on a truly global level we can shed some insight onto appropriate segmentation thresholds by identifying where there are clusterings of people with similar consumption levels.

We identify, from the density curves, a basis for setting such segment thresholds. Imagine a group of people spread around the world but all with broadly similar income per capita in PPP terms (i.e. similar spending or consumption power) then we *might* think of that as a distinct global cluster. If such a group existed, it would be clustered (one would expect with some sort of normal distribution) around an average income point. In a plot of the number of people at each income level we might therefore expect to see a local 'peak' forming with some sort of bell curve centred on this peak. Furthermore, for such a group of people the closer their incomes are (i.e. the less inequality there is within that cluster) the smaller will be the standard deviation of the bell curve (so that the curve will become taller and narrower). And if

the distribution is normal then one standard deviation from the mean would identify the threshold beyond which 15 percent of the distribution would lie. This 85/15 division makes a useful criterion for selecting thresholds.

In Figure 18 we present the density curve for the world in 2010. Consumption per capita (2005 $ PPP) is plotted on a log scale on the horizontal axis. The vertical dashed line (at $730 pa) represents the $2 a day consumption level. The solid line 'population' curve plotted above the horizontal axis represents the number of people living at each consumption level. The area bounded by this 2010 population curve and the horizontal axis represents the entire global population in 2010. That segment of this area that lies to the left of the $2 line represents the proportion of the 2010 population who were living on less than $2 a day (so the ratio of that segment to the entire area of the population curve is the 2010 $2 poverty rate as a percentage of global population). The vertical density axis is dimensionless (it is normalised so that the entire area bounded by the population curve and the horizontal axis aggregates to unity).[17]

The lower curves (plotted negatively) work in the same way but they represent the consumption of the people living at any given level of consumption (as shown on the horizontal axis). The area between the consumption curve and the horizontal axis indicates how much the corresponding population (as indicated by the population curve) collectively consumes per annum (in 2005 PPP $). All the curves are normalised to the global total (population or consumption respectively) in the most recent year of analysis (always 2010 in this paper). Thus, when we plot other population curves in this paper their areas are all relative to, and in proportion to, the

[17] In theory it would be possible to assign a value in terms of actual population count to this axis but that would also require us to specify a bandwidth along the horizontal axis over which that aggregation was calculated. Since this is a log scale, that bandwidth would not translate readily into a simple concept such as 'X thousand people per dollar of consumption'. That approach also allows us to present the population and consumption curves in one graph on the same scale.

2010 global population curve. Similarly consumption curves are all relative to the 2010 global consumption curve. In short, the upper curves show how many people live at each consumption level and the lower curves show how much those people collectively consume. Density curves for the global population are given below (Figures 18 and 19). On these plots we show our three proposed thresholds and also the World Bank's current extreme poverty line ($1.25 a day).

In these curves two main broad 'peaks' can be identified. One is seen in the population curve at low-income levels while the other, less well-defined one, is best seen in the consumption curve. The peaks are more clearly seen in earlier years but when China is excluded can also been seen to persist into the 2000–2010 period. Each of these peaks can be understood as representing the centre of a clustering of individuals. Our segmentation approach identifies a dividing point between the two clusters and then approximately divides each of these clusters into an upper and lower segment.

Figure 18: Global density curve, survey means, filled

Figure 19: Global density curve, survey means, filled – excluding China

We can therefore derive four global consumption 'layers' as follows: first, the 'global absolute poor' – we define this as living under the $2 per capita level. This is not only reasonably close to the midpoint of the low-income peak, it is also well-established as the World Bank moderate international poverty line, which is close to the median poverty line across all developing countries ($2.36 pc in 2008) as well as the regional mean poverty line in sub-Saharan Africa (SSA) and the South Asia Region (SAR) and China, collectively where many of the world's poor live (Ravallion, 2012, p. 25). The

global mean for developing country poverty lines is $4.64 pc which is rather higher than the median because poverty lines can be around $11–$12 pc – the mean in Latin America and the Caribbean and in Eastern Europe and close to $4 – the mean in East Asia and Pacific (Ravallion, 2012, p. 25). In short, $2 pc seems reasonable because it is close both to the global median poverty line and to the mean poverty line in the countries where most of the world's poor live (sub-Saharan Africa, South Asia and China). It is certainly more appropriate than the extreme $1.25 a day poverty line which, as can be seen from the density curves, currently falls consistently below any central point of the low-income peak of the population curve.

Second and third, the 'global insecure' and 'global secure' layers – meaning respectively $2–$10 per capita and $10–$50 per capita. In the density curves these represent the upper half of the lower income population peak and the lower half of the higher income consumption peak respectively. Setting the threshold between these at $10 a day represents a reasonable cut-off point in the overlap between these two peaks. We find that 87 percent of the HIC population lives above $10 pc a day – which fits our 85/15 rule. 98 percent of LIC and LMIC populations are below this level. This threshold therefore broadly separates those living rich-world lifestyles from those living developing world lifestyles. Given that there is an inevitable degree of arbitrariness in the precise location of these thresholds the $10 level seems a reasonable point of separation. If we wished to exactly balance out the LIC/HIC separation we would need a threshold of $7, which would give 94 percent of HIC population above the threshold and 94 percent of LIC and LMIC population below it. Although any line is arbitrary to some extent, $10 has been identified as an approximate 'security' from poverty line or a 'middle class' line in the sense of

security from the risk of falling back into poverty.[18] We therefore propose the $10 threshold and that those living below this level might be referred to as the 'global insecure' segment while those above it would be the 'global secure'.

Fourth, the 'global prosperous' – meaning those living on or above $50/day per capita. This approximates to the midpoint of the higher income peak (about half of HIC consumption is by people living above this level and half is below it). It is also the level below which 87 percent of the HIC population live (so fits a 85/15 'rule'). A further reason for choosing this location as the division between the global secure and the global prosperous is that it coincides with a depression in the consumption peak. Based on the reasoning above that peaks (and subsidiary peaks) represent clusters of individuals with similar preferences, this depression might be understood as the dividing point between two different clusterings within the rich world, perhaps representing a division between two relatively distinct standards of living.[19]

5c. The benefits of global growth since 1990

Who benefited most from global growth since 1990? Figure 20 illustrates how in absolute terms only those in the world's richest 10 percent have seen their

[18] A $10 poverty line was originally proposed by Pritchett (2006). López-Calva & Ortiz-Juarez (2011) refer to a 'vulnerability approach to identifying the middle classes' based on the finding that in Latin America the risk of falling into poverty falls drastically after $10 per capita. Birdsall et al., (2013) noted that $10 is the mean per capita income of those who have completed secondary school across Latin America, suggesting such completion of schooling can be associated with some kind of greater security.

[19] Extending this reasoning, we did also consider introducing a further threshold for the super-rich which would have separated an emerging very high income peak (above the $120 a day level) from the rest of the global prosperous 'layer'. Certainly within this segment there are indications of increasing differentiation along these lines. However, inspection of the underlying data shows that this peak is strongly driven by inequality in the USA. Furthermore, given that there are very substantial errors and exclusions in the measurement in these surveys of the incomes of the super-rich, we decided that separating this peak out would lead to an excessive focus on a trend in the data that is currently not seen across a broad range of countries and, more significantly, that probably is far from representative of the true scale of inequality and consumption at these high income levels (see discussion in Atkinson et al., 2011).

consumption rise by more than the global average, with the main beneficiaries being the global prosperous and particularly those in the global top 1 percent, on family incomes in excess of around $250,000 pa – although caution is needed before drawing conclusions about this, the very richest percentile, because of the known inaccuracies of surveys noted earlier. [20]

Figure 20. Absolute change in income per capita by percentile – world

Global consumption grew by $10 trillion from 1990 to 2010.[21] 15 percent of that global growth accrued to the richest 1 percent of global population (see Tables 2 and 3) while the 3 percent of the global population that are the global prosperous (a group that includes large proportions of the populations of Europe and North America) captured 30 percent of global consumption growth. By contrast, the one-third of the global population that are the now global absolute poor received just 5 percent of the global growth while the two-fifths that are the global insecure received 25 percent of global growth. This is *ex post* and based on the fact there were 35 percent of the global population on $2 or less in 2010. The 5 percent is the share of the total global consumption growth 1990–2010 that went to the bottom 35 percent (i.e. it is the

[20] The $250,000 income figure is derived as follows: the top 1 percent are those with individual consumption levels at or above $75 (in 2005 $ PPP) a day in 2010. If we envisage a typical family of four then that would imply a household consumption of $300 per day reported in direct door-to-door surveys. That is, consumption of around $110,000 a year for the household. If we add, say, 50 percent for income tax and household savings, then this becomes $160,000 a year. We might also want to increase this in line with some portion of the NA/survey ratio (which is typically around 1.6 for high-income countries) to allow for consumption that is not captured in surveys. That would make $75 a day equivalent to a household income of about $250,000 a year.

[21] This $10 trillion figure refers to the change in global aggregate consumption when using survey means and PPP rates. It represents a 79 percent increase in global consumption from 1990 to 2010. Using NA HFC figures without adjustment for the NA/S ratio the growth would be $17 trillion, an 86 percent increase from 1990 to 2010. The average per capita increase on the fractile charts is, of course, lower than this because global population has also risen over the same period.

difference between the aggregate consumption of the bottom 35 percent in 1990 and the bottom 35 percent in 2010).

Table 2: Global population by consumption groups in 2010

Table 3: Shares of global consumption growth

A further perspective is to consider the *ex ante* case. Of the people under $2 in 1990, what share of the $10 trillion did they get? In 1990 about half the world population was on less than $2 a day. Collectively they benefited from less than an eighth of the global growth from 1990 to 2010. In 1990 a little over a third of the world population was on less than $1.25 a day – and collectively they benefited by little more than a twentieth of global growth since then.

If we look at $2 poverty in 2010 the total poverty gap (estimated using survey means) is 6.6 percent of the global growth from 1990 to 2010. In other words, between 1990 and 2010 the one-third of global population who are $2 poor today received 5 percent of global growth while 95 percent went to the rest of the world. An increase in the share of growth to the $2 poor of 7 percent of global growth would have ended $2 poverty. This would have meant that the non-poor would have received 88 percent rather than 95 percent of the global growth in that period. While this is a reduction it is not very dramatic.

Where do the poor, the prosperous and those in-between live? The richest 1 percent are heavily concentrated in North America and Europe where nine-tenths of them live. This includes 15 percent of the US population, 8 percent of the UK population and 2 percent of the entire European Union population. If we turn to the more inclusive global prosperous segment that is 3 percent of the global population

then we find here 36 percent of the US, 14 percent of the UK and 8 percent of the EU population – and 5 percent of the population of Brazil. By contrast, among the global poor and insecure segments we find 90 percent of the Chinese population, 60 percent of Brazil and almost the entire populations of South Asia and sub-Saharan Africa, as opposed to just 12 percent of the US and 13 percent of the EU populations – and only 3 percent of the UK population. The global secure segment includes a fifth of the world's population. Not surprisingly, it includes half the population of the USA and four-fifths of the EU. However, it also includes one-third of Brazil's population and 10 percent of China's, but less than 1 percent of India's population. (see Table 2, Figures 21 to 23).

Figure 21: Estimates of each layer of global population (millions) by region, 1990 and 2010: HFC, survey means, filled

Figure 22: Estimates of each layer of global population (millions) by country income groups, 1990 and 2010: HFC, survey means, filled

Figure 23: Estimates of each layer of global population (millions) for selected countries, 1990 and 2010: HFC, survey means, filled

5d. Relative changes in consumption and patterns of economic growth

The analysis above is based on absolute changes in consumption. Additionally one can look at relative changes in consumption – or by what percentage did consumption levels change for individuals at different locations across the global consumption distribution. We present this data in fractile charts where the horizontal axis represents fractiles rank ordered by level of consumption. The vertical axis represents

the percentage change in consumption per capita from 1990 to 2010. The charts do not include the effects of churn (i.e. that those in the top percentiles today might not be the same as those in the top percentiles in the past) but they do provide a global overview of how, say, the income of 2010's 50th percentile has increased relative to that of 1990's 50th percentile. On each chart we also plot vertical lines at each of the thresholds between our global 'layers' in 2010. When we come to look at regional and country level plots there is a wide variation in the percentage change in per capita consumption so that it is not feasible to plot all these charts to a standard scale. Instead, the vertical scale needs to be varied. In order to provide a common reference point we therefore plot on all the graphs the average global per capita change in HFC, which grew by 47 percent from $3,930 pc pa in 1990 to $5,770 pc pa in 2010. On these graphs, distribution of neutral growth (growth in which everyone saw their income rise in line with the global average) would be represented by a horizontal line at the level of the global average.

Based on the figures we can see that the main beneficiaries of global growth since 1990 have been those in the global insecure segment (between $2 and $10 a day) and below the 70th percentile. They saw their consumption rise by 60 percent or more, considerably above the global average. The main losers were the global secure ($10 to $50 a day), many of whom saw their consumption rise by little more than half of the global average. Among the global prosperous, in general the richer you were the more you benefited, with those in the richest 0.5 percent globally benefiting by more than the global average. Of the 35 percent of the global population that comprise the global absolute poor (below $2 a day) half of them are extremely poor (less than $1.25 a day). Nearly all of the absolute poor benefited at least in proportion to average global growth, but in general for the poorest 50 percent of the world's population the

richer you were, the better you did, so that those around the 50th percentile saw their consumption grow at nearly twice the global average. [22]

The picture, however, looks rather less progressive if China is excluded from the analysis (Figure 24). Most of the global absolute poor (outside China) saw their consumption levels rise by 25 percent to 30 percent and so generally benefited more than the global secure and insecure, most of whom saw rises of around 15 percent. However, the layer that benefited most were the global prosperous, fewer than 5 percent of the global population, who all saw their consumption grow in percentage terms by 30 percent or more (and possibly for the richest by twice that amount), as much or more than any other segment of global society.

Figure 24: Change in consumption 1990 to 2010 at various fractiles, survey means, not filled – world including and excluding China

These fractile charts can also help us identify different types of growth by its 'pro-poor' (or otherwise) nature. There is of course a wealth of literature on such matters and various definitions of 'pro-poor' growth and its predecessors (eg Kakwani et al., 2004; Ravallion, 2004; Ravallion and Chen, 2003) in what has been a long-running discussion on the growth impacts on poverty (see for example, Fosu, 2011; Kraay,

[22] Milanovic (2012a, p. 13) has a similar figure for the period from 1988 to 2008. As does Edward (2006, p. 1681). Milanovic estimated that the top 1 percent globally saw its real income rise by 60 percent in the two decades from 1988. Those between the 75th and 90th percentiles saw a rise of just 10 percent or less (with those around the 80th percentile seeing real falls in income). The main winners, however, were between the 15th and 65th percentiles where real income rose by between 60 percent and 80 percent. We find very broadly similar trends when looking at the period from 1990 to 2010 and present these here and below. We did also look at the 1988 to 2008 period that Milanovic used and our findings there were similar to those we describe here for 1990 to 2010. The main differences to Milanovic's results are that in our results the main winners (other than the richest few percent) appear to be between the 45th and 75th percentiles whose incomes rose by between 60 percent and (almost) 90 percent – a smaller group than Milanovic found. And while the main losers were in the region of the 80th to 90th percentiles, even the worst losers still saw a small real consumption rise over the two decades.

2006; Loayza and Raddatz, 2010) and more recently the evolution of debates into 'inclusive growth' (see for discussion, Klasen, 2010; Ranieri and Ramos, 2013). Here we are making an observation about a small number of stylized patterns relating to who benefits most from growth as per the growth incidence curves. For example, in the growth incidence curves, 'pro-poor growth', in which the poorer end of the distribution benefits disproportionately relative to the rich, with the overall effect therefore of reducing inequality, would show as a line sloping down from upper-left to lower-right on these graphs. And by the same reasoning, growth that increases inequality and benefits the rich proportionally more than the poor would slope up from lower-left to upper-right. With this in mind, we can look at how growth and inequality have interacted in the same period in different countries and regions. The USA (Figure 25) is an example of consistently pro-rich growth across all consumption levels. In contrast, in the EU, growth has generally been relatively evenly spread across 90 percent of the population, with the poorest 10 percent, who find themselves in the global insecure layer, being left behind.

In Latin America, growth appears at first sight to have been reasonably distributionally neutral (compared say to the USA) (Figure 26) but this is largely accounted for by growth in Brazil which has particularly benefited the half of Brazil's population that are in the global insecure segment. When Brazil is removed from the analysis, Latin American growth looks much more pro-rich with most of the consumption increases accruing to those on more than $10 a day. In the first instance this finding might appear to be counter to the finding of falling inequality across Latin America in practically all countries since the 2000s and that in many countries the decline was sharper than the increase in the 1990s (see Lustig et al., 2013). However, as the annex table A1 shows in 1990, Latin America minus Brazil had a

lower overall Gini, both total and within country, compared to 1985 or 1995. Thus one could posit that for LAC (both with and without Brazil) the lower gini in 1990 tends to obscure that rises in total and within country ginis from 1995 to 2000 were more than reversed from 2000 to 2010 so that the current trends in Latin America might be less pro-rich.

In contrast to Latin America, in East Asia, growth has been relatively distribution neutral with all segments benefiting more or less equally (see Figure 27).[23] China's growth has been 'pro-rich', in the sense that the benefits are highest at the top of the distribution but even the lowest incomes have seen very large increases in per capita consumption (Figure 28). In a similar vein, one could say India's growth is also 'pro-rich' (Figure 29). The top 20 percent have seen most of the benefits, and the top 5 percent may have done particularly well. Across the rest of the population, who are almost all in the global poor or global insecure segments, growth has nevertheless been mainly distribution neutral. In Nigeria, growth is similarly pro-rich but it is also both anti-poor and anti-middle too – meaning, surprisingly, the benefits of growth are actually negative at the lower end and concentrated at the top of the distribution (see Figure 29). In addition to the pro-rich growth pattern, four further stylised patterns are discernable from country level analysis: first, there are examples of 'pro-poor growth', by which we mean those at the bottom end of the distribution have benefited most – examples include Ethiopia (see Figure 28), South Africa, Malawi and Mali. Second, there are examples of what one might call 'pro-middle growth' where those in the middle have benefited most – examples include Brazil (Figure 29). Third, there is 'anti-middle growth' such as in Zambia (Figure 29).

[23] When looking at these curves one should focus on overall trends and differences across the curves and be especially wary of reading too much significance into large changes at the extreme ends (typically the last 3 percent of the distribution) of the graphs – as occurs, for example, in the EAP without China curve above the $50 a day level. At these ends the analysis can be particularly sensitive to slight changes in the data so that these large changes may not be reliable representations of actuality.

Finally, there is 'equitable growth', meaning equitable across the distribution– this is the case in Vietnam (Figure 28), Nepal and the Philippines. These findings suggest a wider typology of patterns of growth that might be more useful than a focus on single-coefficient measures, such as poverty headcounts or Gini or Theil, when assessing the distribution of economic growth.

Figure 25. USA and EU

Figure 26. LAC with and without Brazil

Figure 27. EAP with and without China

Figure 28. China, Vietnam, Ethiopia and Indonesia

Figure 29. Brazil, India, Nigeria and Zambia

6. CONCLUSIONS

In this paper we have provided new estimates of the evolution of inequality between and within countries. In conclusion, we can note three points:

First, it may come as a surprise, but we find that total global inequality has been relatively static for most of the period since the late 1980s. This is because a slow but steady rise in within-country inequality was largely, but not completely, offset by a gradual decline in between-country inequality. Since 2005 between-country inequality has been falling more quickly than before and as a result total global inequality has also started to fall, quite quickly in the last few years.

Not surprisingly, but little noted, is the 'China effect' or the role of China in determining these trends. Indeed, the picture looks rather different when China is excluded: in the rest of the world outside China between-country inequality rose in

the 1980s and 1990s but has then stayed relatively constant since 2000. Throughout this entire period within-country inequality has overall been remarkably constant – as some countries have become less equal, others have become more so. In short, in the last 20 to 30 years, falls in total global inequality, and in global between-country inequality, and rises in global within-country inequality are all predominantly attributable to rising prosperity in China.

Second, we argue that global society can be divided into four global 'layers' of the 'global absolute poor' (<$2 a day); the 'global prosperous' (>$50 a day) and those in-between, specifically, the 'global insecure' ($2 to $10 a day) and the 'global secure' ($10 to $50 a day). One might also add a top layer of the 1 percent of population on over $75 a day (equivalent to a family income of around $250,000 pa in 2005 $PPP). At a global level, we still see a global structure of two relatively homogeneous clusters (the poor/insecure and secure/prosperous clusters) but we also detect the emergence of a rapidly changing and heterogeneous 'new global middle' that is filling the space of the missing middle in this otherwise remarkably stable and binary (twin-peak or dumb-bell) global consumption distribution. However, most of the 'structural' change in the distribution of global consumption is confined to the upper middle-income countries (UMICs), notably China. If we omit the UMICs, then the long-standing twin-peak 'dumb-bell' of a bipolar world of two highly differentiated clusters persists largely unchanged. This leads us to question whether the emerging global middle really does represent an evolution away from the historical twin-peak world or whether it simply represents a transition phase as some elements in emerging economies move from the poor/insecure cluster into the secure/prosperous cluster.

Third, we find that global consumption grew by $10 trillion from 1990 to 2010 (based on survey means and PPP rates), of which we estimate 15 percent went to the richest 1 percent of the global population. At the other end of the distribution, collectively, those under $2 in 1990 (53 percent of the population at that time) benefited from less than an eighth of the global growth from 1990 to 2010; and the 37 percent of the world population on less than $1.25 a day in 1990 collectively benefited by little more than a twentieth of global growth.

ACKNOWLEDGEMENTS

We would like to thank the following for comments on earlier drafts: Nora Lustig, Stephanie Seguino, Borge Wietzke, Martin Evans, Simon Maxwell, Owen Barder and Richard Manning.

REFERENCES

Atkinson, A., Piketty, T., Saez, E., 2011. Top incomes in the long run of history. Journal of Economic Literature, 49(1), 3–71.

Banerjee, A., Duflo. E., 2008. What Is Middle Class about the Middle Classes around the World? Journal of Economic Perspectives, 22(2), 3-28.

Bhalla, S., 2002. Imagine There's no Country: Poverty, Inequality and Growth in the Era of Globalization. Institute for International Economics, Washington, DC.

Bhalla, S., 2007. Second Among Equals: The Middle Class Kingdoms of India and China. Peterson Institute for International Economics: Washington DC.

Birdsall, N., Lustig, N., Meyer, C., forthcoming. The New Poor in Latin America. Centre for Global Development (CGD) Working Paper. CGD, Washington, DC.

Bourguignon, F., Morrisson, C., 2002. Inequality among world citizens: 1820–1992. American Economic Review, 92(4), 727–744.

Carroll, W.K., 2012. Global, transnational, regional, national: the need for nuance in theorizing global capitalism. Critical Sociology, 38(3), 365–371.

Cowell, F., 2000. Measuring Income Inequality. Harvester Wheatsheaf, London.

Dikhanov, Y., Ward, M., 2002. Evolution of the Global Distribution of Income in 1970–99. Proceedings of the Global Poverty Workshop. Initiative for Policy Dialogue, New York.

Dowrick, S., Akmal, M., 2005. Contradictory trends in global income inequality: a tale of two biases. Review of Income and Wealth, 51(2), 201–229.

Edward, P., 2006. Examining inequality: who really benefits from global growth? World Development, 34(10), 1667–1695.

Edward, P., Sumner, A., 2013. The Future of Global Poverty in a Multi-Speed World. Centre for Global Development (CGD) Working Paper. CGD, Washington, DC.

Fosu, A., 2011. Growth, Inequality and Poverty Reduction in Developing Countries: Recent Global Evidence. UNU WIDER Working Paper. UNU-WIDER, Helsinki.

Goda, T., 2013. Changes in Income Inequality from a Global Perspective: An Overview. Post-Keynesian Economics Study Group Working Paper 1303. PKSG, Cambridge.

Goldthorpe, J.H., McKnight, A., 2006. The economic basis of social class, in: Morgan, S.L., Grusky, D.B., Fields, G.S. (Eds.), Mobility and Inequality: Frontiers of Research from Sociology and Economics. Stanford University Press, Stanford, CA, pp. 109–136.

Hillebrand, E., 2008. The global distribution of income in 2050. World Development, 36(5), 727–740.

Kakwani, N., Khandker, S., and Son, H., 2004. Pro-Poor Growth:

Concepts and Measurements with Country Case Studies". Working Paper Number 1 International Poverty Centre, Brasilia.

Kharas, H., 2010a. The Emerging Middle Class in Developing Countries. OECD Development Centre. OECD, Paris.

Klasen, S., 2010. Measuring and Monitoring Inclusive Growth: Multiple Definitions, Open Questions, and Some Constructive Proposals, ADB Sustainable Development Working Paper Series, No. 12. Manila, Asian Development Bank.

Kraay, A., 2004. When is Growth Pro-Poor? Cross-Country Evidence. World

Bank Working Paper Number 3225. World Bank: Washington, DC.

López-Calva, L.F., Ortiz-Juarez, E., 2011. A Vulnerability Approach to the Definition of the Middle Class. Policy Research Working Paper 5902. World Bank, Washington, DC.

Linder, S.B., 1961. An Essay on Trade and Transformation. Wiley, New York.

Loayza, N., Raddatz, C., 2010. The Composition of Growth Matters for Poverty Alleviation. Journal of Development Economics, 93, 137–151.

Lustig, N., Lopez-Calva, L., Ortiz-Juarez, E., 2011. The Decline in Inequality in Latin America: How Much, Since When and Why. Tulane Economics Working Paper 1118. Tulane University, New Orleans, LA.

Lustig, N., Lopez-Calva, L., Ortiz-Juarez, E., 2013. Declining inequality in Latin America in the 2000s: the cases of Argentina, Brazil, and Mexico. World Development, 44, 129–141.

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

Milanovic, B., 2005. Worlds Apart: Measuring Global and International Inequality. Princeton University Press, Princeton, NJ.

Milanovic, B., 2011. Global Inequality: From Class to Location, from Proletarians to Migrants. World Bank Working Paper 5820. World Bank, Washington, DC.

Milanovic, B., 2012a. Global Income Inequality by the Numbers: In History and Now – an Overview. World Bank Working Paper 6259. World Bank, Washington, DC.

Milanovic, B., 2012b. Global inequality recalculated and updated: the effect of new PPP estimates on global inequality and 2005 estimates. Journal of Economic Inequality, 10(1), 1–18.

Milanovic, B., Yitzhaki, S., 2002. Decomposing world income distribution: does the world have a middle class? Review of Income and Wealth, 48(2), 155–178.

Nel, P., 2012. The Politics of Economic Inequality. Palgrave MacMillan, Basingstoke.

Palma, G., 2011. Homogeneous middles vs. heterogeneous tails, and the end of the 'Inverted-U': the share of the rich is what it's all about. Development and Change, 42(1), 87–153.

Pritchett, L., 2006. Who is not Poor? Proposing a Higher International Standard for Poverty. World Bank Research Observer. 21(1): 1-23

Quah, D., 1996. Twin peaks: growth and convergence in models of distribution dynamics. Economic Journal, 106 (437), 1045–55.

Ranieri, R., Ramos, R., 2013. Inclusive Growth: the Building up of a Concept, Working Paper No. 104. Brasília, International Policy Centre for Inclusive Growth.

Ravallion, M., 2004. Pro-Poor Growth: A Primer. Policy Research Working Paper, Working Paper 3242, March, Washington DC: World Bank.

Ravallion, M., 2010. The developing world's bulging (but vulnerable) middle class. World Development, 38(4), 445–454.

Ravallion, M., Chen, C., 2003. Measuring Pro-Poor Growth. Economic Letters, Vol. 78, pp. 93-99.

Ravallion, M., Chen, S., 2012. Monitoring Inequality. World Bank, Washington, DC.

Robinson, W.I., 2012. Global capitalism theory and the emergence of transnational elites. Critical Sociology, 38(3), 349–363.

Sala-i-Martin, X., 2002. The World Distribution of Income (Estimated from Individual Country Distributions). NBER Working Paper 8933. NBER, Cambridge, MA.

Sala-i-Martin, X., 2006. The world distribution of income: falling poverty and ... convergence, period. Quarterly Journal of Economics, 71(2), 351–397.

Standing, G., 2011. The Precariat: The New Dangerous Class. Bloomsbury Academic, London.

Sudhir, A., Segal, P., 2008. What do we know about global income inequality? Journal of Economic Literature, 46(1), 57–94.

Sumner, A., 2012. Poverty, vulnerability and class: the expanding non-polar groups and development. Public Administration and Development, 32(4–5), 444–454.

Torche, F., López-Calva, L.F., forthcoming. Stability and vulnerability of the Latin American middle class, in: Katherine Newman (Ed.), Dilemmas of the Middle Class around the World. Oxford University Press, New York.

UNU WIDER. 2008. World Income Inequality Database. WIDER, Helsinki, Finland.

Wade, R., 2004. Is globalization reducing poverty and inequality? World Development, 32(4), 567–589.

World Bank, 2012. World Development Indicators 2012. World Bank, Washington, DC

Figures

Figure 1: Global Gini coefficient, survey means (not filled)

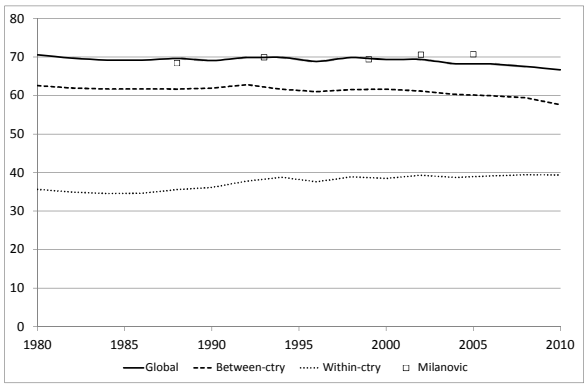

Figure 2: Global Gini coefficient excluding China, survey means (not filled)

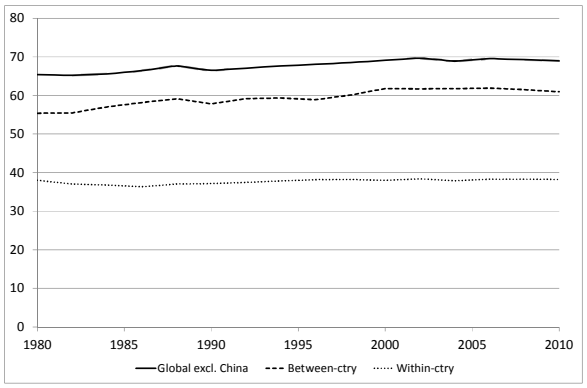

Figure 3: Global Theil Index, survey means (not filled)

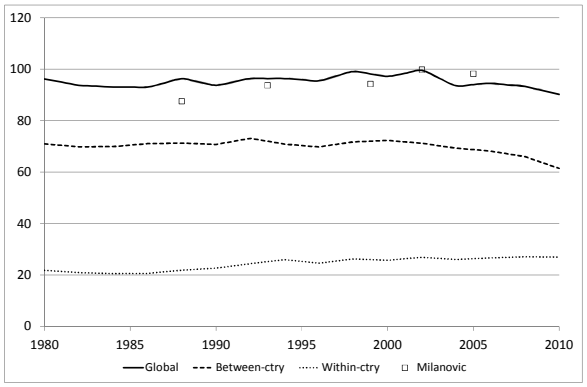

Figure 4: Global Theil Index excluding China, Theil Index, survey means (not filled)

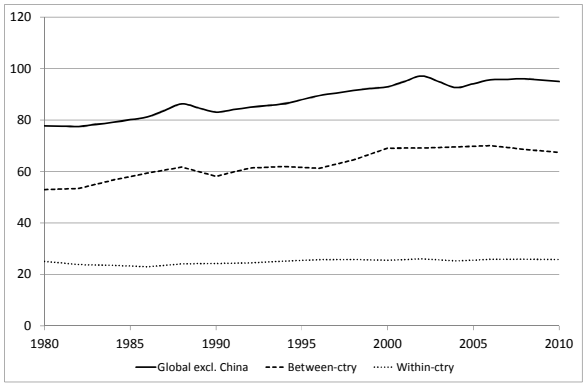

Figure 5: Within-country Theil component as percentage of global Theil Index, survey means (not filled)

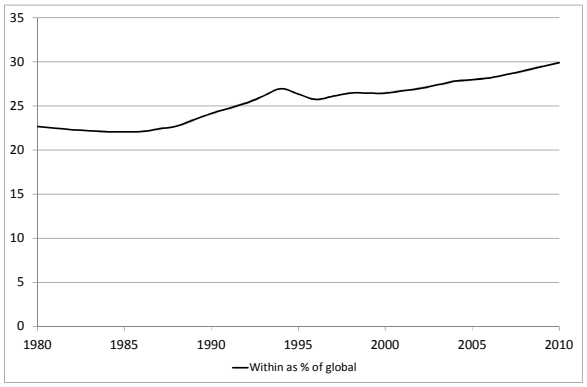

Figure 6: World excl. China, Within-country Theil component as percentage of global Theil Index, survey means (not filled)

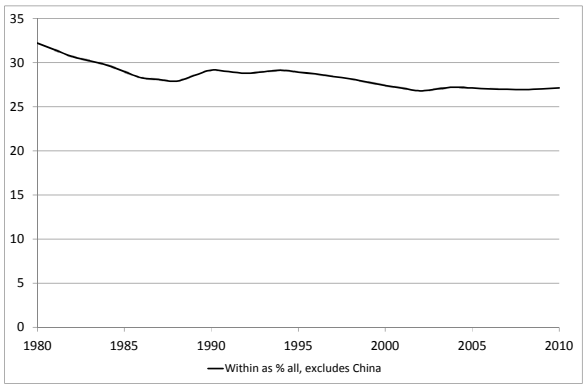

Figure 7: The 'twin-peak' global distribution of the mid-1980s

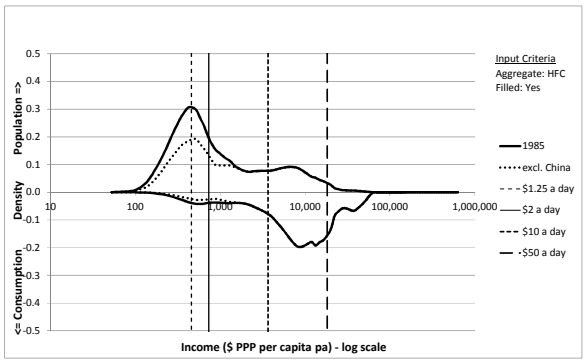

Figure 8: The 2010 global distribution and loss of the 'twin peaks' or 'dumb-bell'

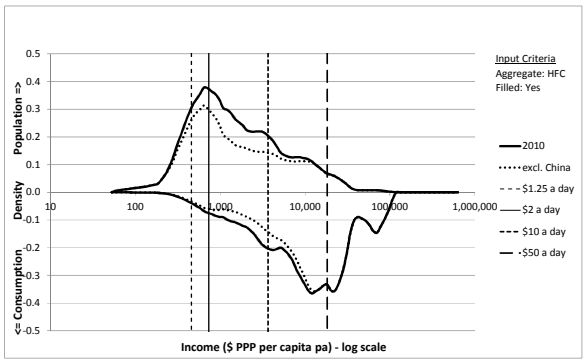

Figure 9: Density curves for current HICs

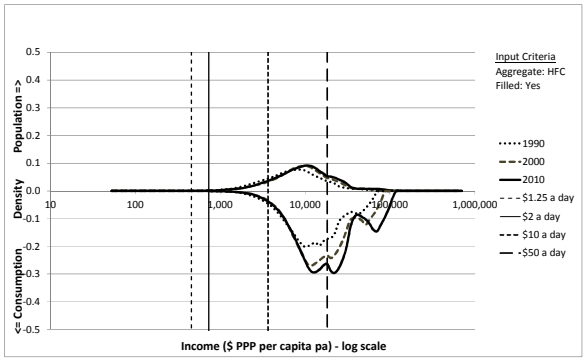

Figure 10: Density curves for current LICs and LMICs

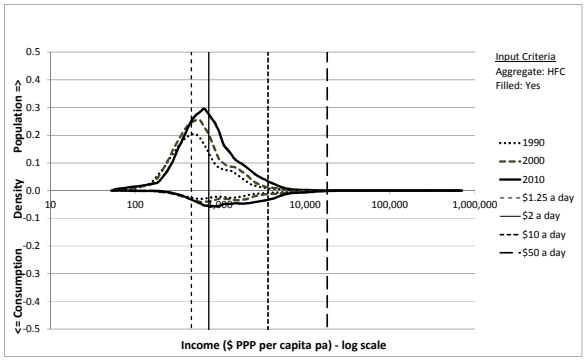

Figure 11: Density curves for current UMICs

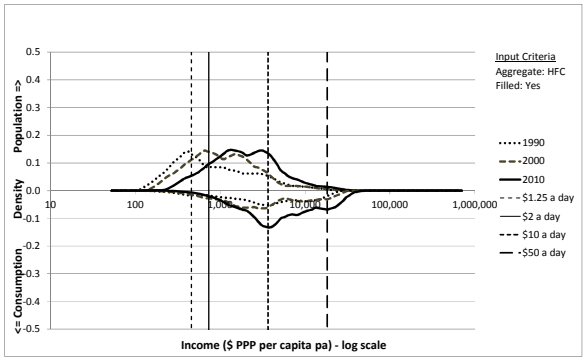

Figure 12: Density curves for current LICs and LMICs excluding India

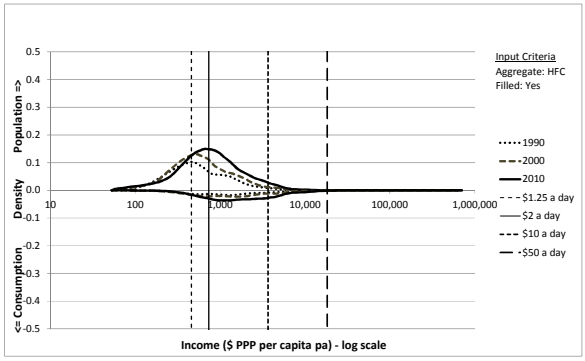

Figure 13: Density curves for current UMICs excluding China

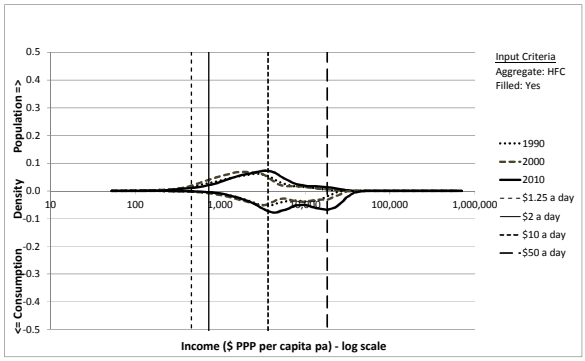

Figure 14: Density curve for India

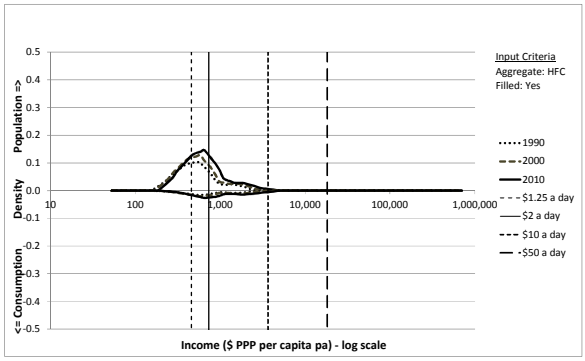

Figure 15: Density curve for China

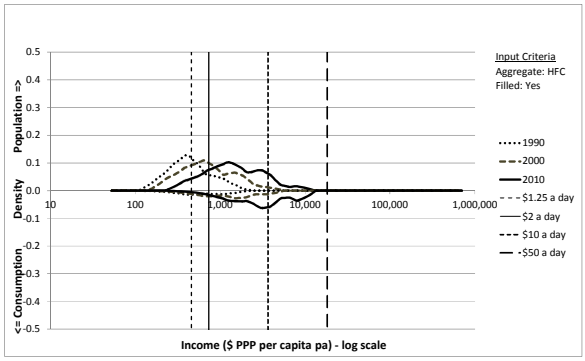

Figure 16: Density curves for all countries excluding current UMICs

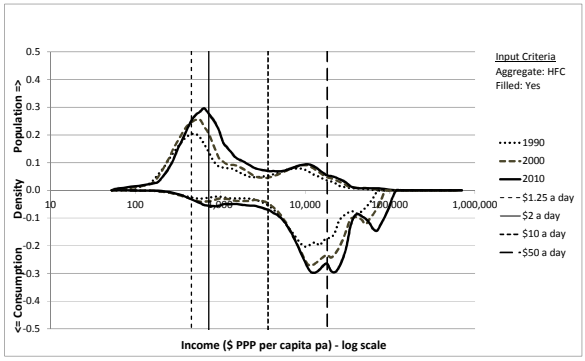

Figure 17. Gini coefficient, 1980–2010 in selected populous countries

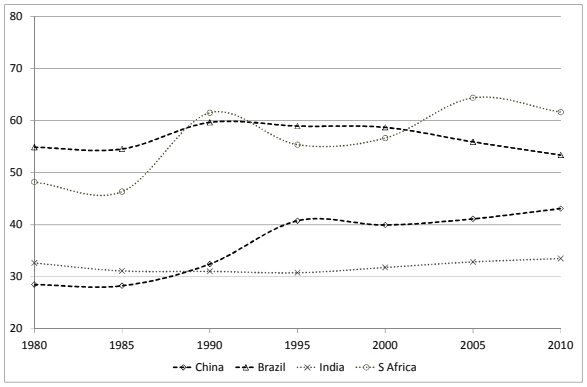

Figure 18: Global density curve, survey means, filled

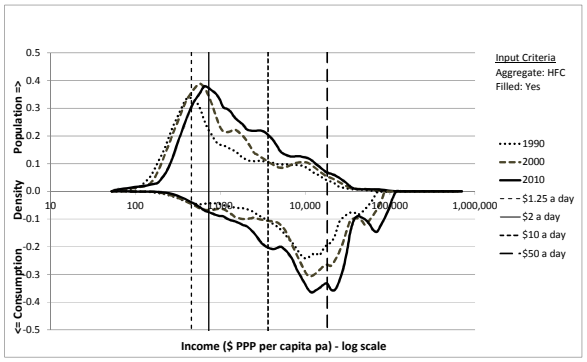

Figure 19: Global density curve, survey means, filled – excluding China

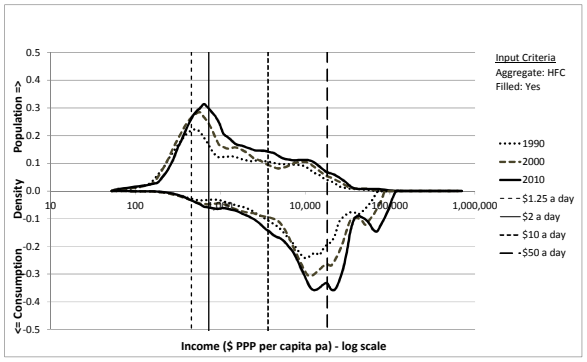

Figure 20. Absolute change in income per capita by percentile – world

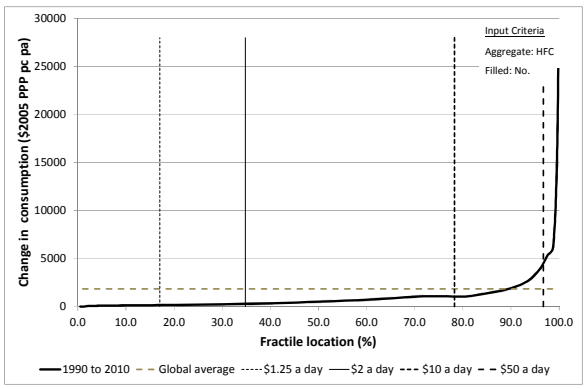

Figure 21: Estimates of each layer of global population (millions) by region, 1990 and 2010: HFC, survey means, filled

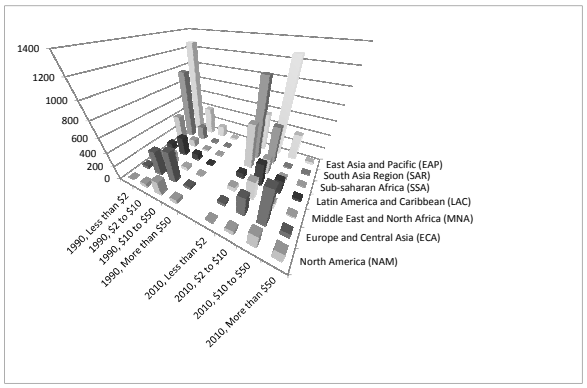

Figure 22: Estimates of each layer of global population (millions) by country income groups, 1990 and 2010: HFC, survey means, filled

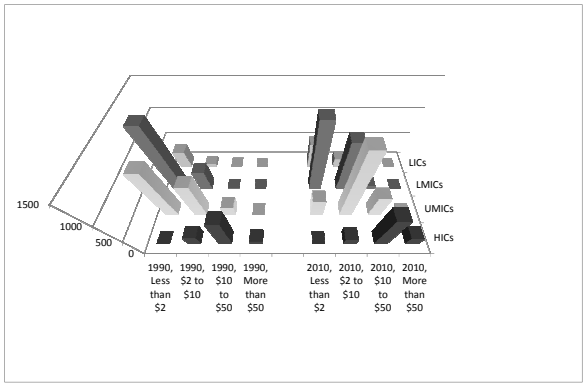

Figure 23: Estimates of each layer of global population (millions) for selected countries, 1990 and 2010: HFC, survey means, filled

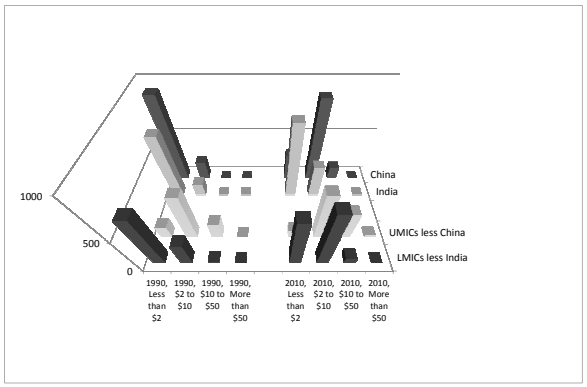

Figure 24: Change in consumption 1990 to 2010 at various fractiles, survey means, not filled – world including and excluding China

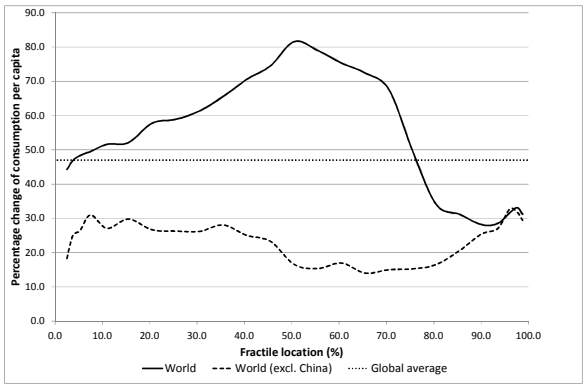

Figure 25. USA and EU

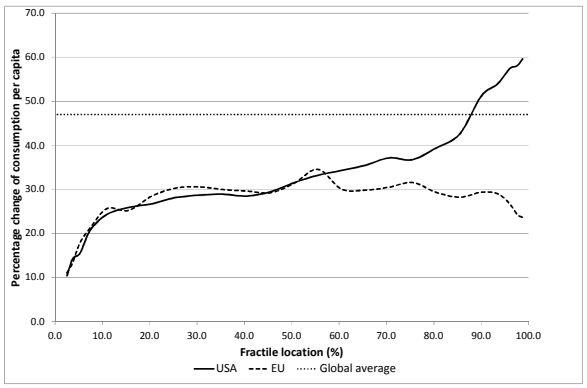

Figure 26. LAC with and without Brazil

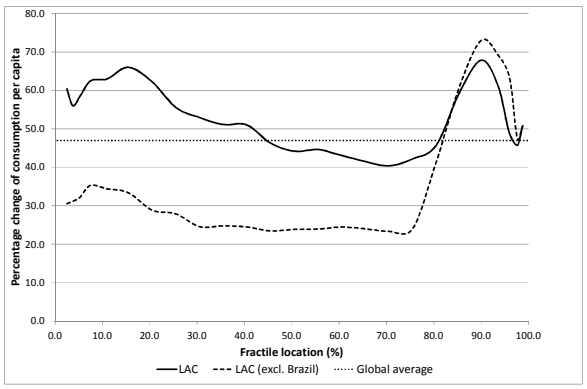

Figure 27. EAP with and without China

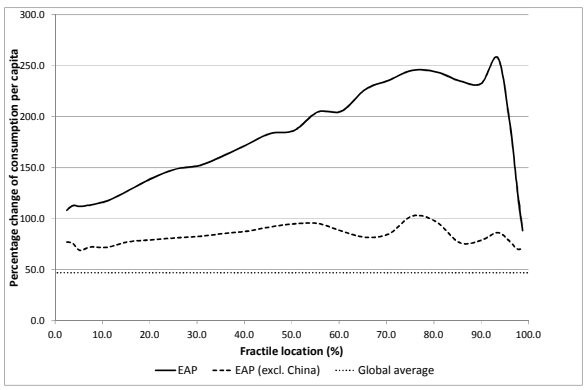

Figure 28. China, Vietnam, Ethiopia and Indonesia

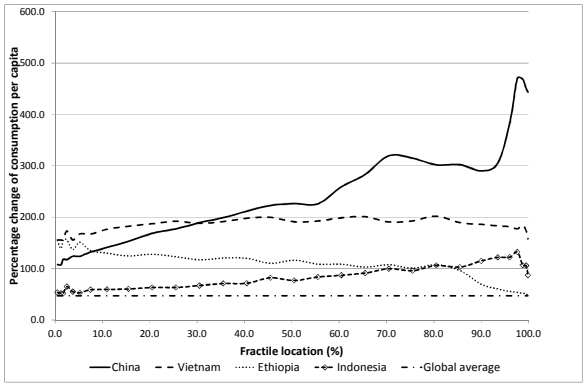

Figure 29. Brazil, India, Nigeria and Zambia

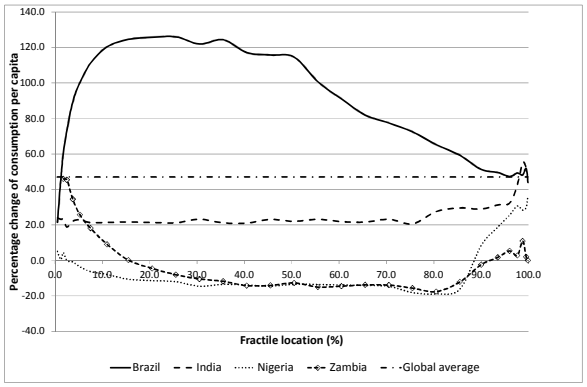

Table 1: Coverage (cov.) of analysis and effects of estimating HFC and filling distributions

	Source data coverage			After estimating missing HFC			After filling missing distributions		
Year	No. of countries	Pop'n cov. (%)	Consumption cov. (%)	No. of countries	Pop'n cov. (%)	Consumption cov. (%)	No. of countries	Pop'n cov. (%)	Consumption cov. (%)
1980	62	71.7	72.6	79	81.2	83.9	132	85.9	87.7
1990	97	84.4	81.0	131	94.0	92.6	167	96.3	94.3
2000	118	87.2	82.7	156	96.2	91.2	181	97.4	92.5
2010	102	83.4	78.4	135	91.9	80.1	178	96.6	89.6

Source: Authors' estimates based on GrIP v1.0. Percentages are of global totals.

Table 2: Global population by consumption groups in 2010

	Less than $2	$2-$10	$10-$50	$50+	$75+	Global total
	Global Absolute Poor	Global Insecure	Global Secure	Global Prosperous	Top 1%	
Total (millions)	2407	2910	1351	227	69	6894
As % of global population	35	42	20	3	1	100
Regional distribution (millions)						
East Asia and Pacific (EAP)	542	1267	370	23	3	2202
Europe and Central Asia (ECA)	27	269	542	54	11	891
Latin America and Caribbean (LAC)	70	317	181	20	3	589
Middle East and North Africa (MNA)	48	262	68	5	2	383
North America (NAM)	0	43	175	125	50	344
South Asia Region (SAR)	1092	537	4	0	0	1633
sub-Saharan Africa (SSA)	627	214	12	0	0	854
Distribution by income category (millions)						
LICs	543	177	1	0	0	722
LMICs	1459	1096	70	0	0	2625
UMICs	403	1488	486	28	3	2405
HICs	1	148	793	199	66	1142

Source: Authors' estimates based on GrIP v1.0. Note: Numbers are derived using filled analysis so as to most closely match global aggregates.

Table 3: Shares of global consumption growth

Global segment (in 2010 unless noted)	Share of global population (%)	Share of global consumption growth 1990 to 2010 (%)
Global Absolute Poor (<$2)	34.9	5.1
Global Insecure ($2.01-$10)	42.3	24.7
Global Secure ($10.01-$50)	19.5	41.4
Global Prosperous ($50.01+)	3.3	28.7
Top 1% ($75+)	1.0	14.9
The $1.25 poor in 1990	36.8	5.7
The $2 poor in 1990	53.1	11.7
The $1.25 poor in 2010	18.2	1.8

Source: Authors' estimates based on GrIP v1.0.

Annex table

Table A1: Selected regional Ginis (Option 1, HFC, No fill. Regional groupings exclude HICs)

	Total Gini coefficient							Within country Gini component						
	1980	**1985**	**1990**	**1995**	**2000**	**2005**	**2010**	**1980**	**1985**	**1990**	**1995**	**2000**	**2005**	**2010**
World	70.6	69.2	69.1	69.4	69.4	68.2	66.7	35.6	34.6	36.1	38.2	38.5	38.9	39.4
World, excl. China	65.4	66.0	66.5	67.9	69.1	69.2	68.9	38.0	36.6	37.2	38.0	38.0	38.1	38.2
EAP	41.8	35.6	39.0	44.7	42.2	42.4	44.5	30.6	30.0	33.4	39.9	39.2	40.4	41.7
LAC	53.1	51.5	51.9	55.5	56.3	54.4	52.7	50.8	49.6	50.6	53.1	54.4	52.9	50.3
SAR	32.9	31.4	31.6	31.4	32.7	33.6	33.8	32.3	30.8	31.0	30.9	32.1	32.9	32.9
SSA	56.9	53.7	56.7	53.6	52.4	54.5	55.1	48.7	42.9	47.2	44.8	43.3	43.6	43.4
LICs	30.8	41.6	49.1	42.3	41.8	41.2	43.0	30.3	36.2	40.3	38.8	38.4	38.3	37.0
LMICs (excl. India)	52.8	46.3	46.0	47.6	45.1	44.6	45.4	39.8	35.8	35.8	37.7	37.6	38.2	37.9
UMICs (excl. China)	51.2	49.7	47.0	51.0	50.2	48.7	48.7	48.3	46.6	44.6	47.7	47.5	47.1	46.2
LICs and MICs (excl. India and China)	61.0	58.8	57.6	58.4	56.4	56.4	58.6	43.0	40.7	40.5	42.2	41.8	41.6	40.6
HICs	39.2	38.5	40.4	41.9	39.3	39.3	42.5	34.9	33.5	35.0	36.2	33.7	33.4	36.2
EAP less China	54.4	47.6	47.8	48.8	47.2	45.2	46.7	37.6	35.7	36.0	36.9	37.0	38.4	37.8
LAC less Brazil	51.8	49.6	46.5	52.6	54.4	53.3	51.9	48.2	46.5	44.6	49.7	51.9	51.2	48.6
SAR less India	33.5	32.1	32.8	33.2	35.2	35.2	34.1	30.8	29.7	30.8	31.4	33.2	32.9	31.1
SSA less South Africa	53.0	47.3	50.4	49.1	48.1	46.7	48.6	48.7	42.4	45.5	43.7	42.1	41.7	41.9
HICs less USA	35.5	32.4	35.5	36.6	34.1	33.4	32.7	33.0	30.0	32.0	32.7	31.2	30.6	29.9
China	28.5	28.3	32.4	40.8	39.9	41.1	43.1							
Brazil	54.9	54.5	59.6	58.9	58.7	55.9	53.4							
India	32.7	31.1	31.0	30.8	31.8	32.9	33.5							
South Africa	48.2	46.3	61.5	55.4	56.6	64.4	61.6							
USA	39.7	41.5	42.5	44.5	39.6	39.6	46.1							
EU	34.0	30.0	34.0	39.6	37.9	35.9	34.3							