Working Paper

THE POOR, THE PROSPEROUS AND THE 'INBETWEENERS':

A FRESH PERSPECTIVE ON GLOBAL SOCIETY, INEQUALITY AND GROWTH

Working Paper number 122 March, 2014

Peter Edward
Newcastle University Business School

Andy Sumner
International Development Institute,
King's College London

International Policy Centre for Inclusive Growth

International Policy Centre for Inclusive Growth (IPC - IG)
United Nations Development Programme

SBS, Quadra 1, Bloco J, Ed. BNDES, 13º andar
70076-900 Brasilia, DF - Brazil
Telephone: +55 61 21055000

E-mail: ipc@ipc-undp.org ▪ URL: www.ipc-undp.org

The International Policy Centre for Inclusive Growth is jointly supported by the United Nations Development Programme and the Government of Brazil.

Working Papers are available online at www.ipc-undp.org and subscriptions can be requested by email to ipc@ipc-undp.org

Print ISSN: 1812-108X

THE POOR, THE PROSPEROUS AND THE 'INBETWEENERS':

A FRESH PERSPECTIVE ON GLOBAL SOCIETY, INEQUALITY AND GROWTH

Peter Edward and Andy Sumner*

What has happened to inequality between and within countries since 1990? In this paper we explore who have been the winners and losers from global growth since 1990. We find that falls in total global inequality in the last 30 years are predominantly attributable to rising prosperity in China. We also identify a persistent global structure of two relatively homogeneous clusters (the poor/insecure and secure/prosperous). We detect the emergence of a 'new global middle' but question whether this implies the end of the historical two-cluster world rather than merely a transition as some people move from the poor/insecure cluster into the secure/prosperous cluster. Nevertheless, we do identify five different stylised patterns of national growth: pro-poor growth (e.g. Ethiopia); pro-middle growth (e.g. Brazil); anti-poor growth (e.g. Nigeria); anti-middle growth (e.g. Zambia) and equitable growth (e.g. Vietnam). We also find that 15 per cent of growth from 1990 to 2010 went to the world's richest 1 per cent, while just a modest amount of redistribution would have ended $2 poverty. If the share of global growth between 1990 and 2010 flowing to those who were living on under $2/day in 2010 had increased from 5 per cent to just 12 per cent, this would have been sufficient to end $2 poverty today. Persistence of global poverty, it seems, is not due to insufficient global growth but to a reluctance among the secure/prosperous cluster to forego a small share of their benefits from global growth in favour of fairly modest redistribution to the global poor.

Key Words: Poverty; Inequality; Economic Development

JEL Codes: D63; I32

* Respectively, Newcastle University Business School and King's International Development Institute, King's College London. Correspondence to peter.edward@ncl.ac.uk and andrew.sumner@kcl.ac.uk. Many thanks for important comments to Simon Maxwell, Owen Barder, Richard Manning, Martin Evans, Borge Wietzke and Stephanie Seguino.

EXECUTIVE SUMMARY

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 since 1990. We find that total global inequality was relatively static from the late 1980s to 2005, with rising inequality within countries largely offset by falling inequality between countries. Since 2005 between-country inequality has been falling more quickly than before, and as a result total global inequality has also fallen. We find that falls in total global inequality and global between-country inequality and rises in global within-country inequality in the last 20 to 30 years are all predominantly attributable to rising prosperity in China. The picture looks rather different when China is excluded. Throughout this entire period within-country inequality in the rest of the world has overall been remarkably constant— as some countries have become less equal, others have become more so —while between-country inequality has actually increased slightly.

We identify a persistent global structure of two relatively homogeneous clusters (the poor/insecure and secure/prosperous), but we also detect the emergence of a rapidly changing and heterogeneous 'new global middle'. However, most of the 'structural' change in the distribution of global consumption is confined to the upper middle-income countries (UMICs). This leads us to question whether the emerging global middle really does represent an evolution away from the historical two-cluster world or simply represents a transition phase as some elements in emerging economies move from the poor/insecure cluster into the secure/prosperous cluster. Nevertheless, we do identify five different stylised patterns of national growth as follows: pro-poor growth (e.g. Ethiopia); pro-middle growth (e.g. Brazil); anti-poor growth (e.g. Nigeria); anti-middle growth (e.g. Zambia) and equitable growth (e.g. Vietnam). We also find that 15 per cent of global consumption growth from 1990 to 2010 went to the richest 1 per cent of the global population. At the other end of the distribution, the 53 per cent of the population living on under $2/day in 1990 benefited from less than an eighth of that global growth, and the 37 per cent living on less than $1.25/day benefited by little more than a 20th of that growth. A modest amount of redistribution would have ended $2 poverty: if the share of global growth since 1990 flowing to the 35 per cent of the global population who were living on $2/day in 2010 had increased from 5 per cent to 12 per cent, this would have been sufficient to end poverty at $2/day. In short, the persistence of global poverty seems to have little to do with there being insufficient global growth and a lot more to do with a lack of collective will among the secure/prosperous cluster to forego a small share of their benefits from global growth in favour of a fairly modest amount of redistribution to the global poor.

1 INTRODUCTION

The interplay of inequality between and within countries, the relative contribution of each to overall global inequality and the implications this has for who benefits (and by how much) from recent global growth has become a significant avenue for research and analysis.[1] Indeed, a number of recent empirical papers have discussed changes in inequality, some with reference to 'class' in the contemporary period of globalisation, typically considered to be since the collapse of the Soviet Union and the end of the Cold War in the late 1980s (some on inequality —e.g. Milanovic, 2011; 2012; Palma, 2011; Ravallion and Chen, 2012 —and others on the emergence of 'new middle classes' —e.g. Kharas, 2010a; 2010b; Ravallion, 2010).

In this paper we provide new estimates of the evolution of inequality between and within countries and then focus on inequality in a multi-layered world. These are 'new' in the sense they are derived from a purpose-built model. We use globally standardised absolute consumption thresholds to identify and consider the fortunes of four global 'layers' as follows: the 'global absolute poor' (<$2/day); the 'global prosperous' (>$50/day) and those in between —specifically, the 'global insecure' ($2–10/day) and the 'global secure' ($10–50/day). We prefer to call these global 'layers', rather than 'classes', because the conflation of differences in per capita consumption levels with 'class' is —of course —problematic, since class is a social and political identity that does not automatically follow, even at the national level let alone the global level, from consumption (or income) level (see discussion in Sumner, 2012).

We derived estimates using the GrIP ('Gr'owth, 'I'nequality and 'P'overty) model (version 1.0), which has been developed from an earlier model described in Edward (2006) and is discussed briefly below and in more depth in Edward and Sumner (2013). The paper itself is structured as follows: Section 2 reviews recent empirical studies. Section 3 outlines the GrIP model. Section 4 asks what has happened to inequality between and within countries over the last 20 years. Section 5 then presents the 'layers' approach to inequality and estimates of the distribution of benefits during the period of globalisation since the end of the Cold War up to 2010. Section 6 concludes.

2 REVIEW OF STUDIES

2.1 POINTS OF DEPARTURE

A review of peer-reviewed studies published since 2000 relating to the empirical study of long-term trends in income/consumption inequality identified more than 70 relevant papers in economics and development journals.[2] A parallel review of the distribution of the benefits of economic growth produced more than 50 further papers.[3] These papers point towards a high level of interest and concern in the ways that growth, inequality and poverty interact. However, because of the purchasing power parity (PPP) revisions in 2008 (for 2005 PPPs), many of these studies have had a relatively short 'shelf-life'.

Several recent papers are worthy of note because they are especially relevant to the analysis later in this paper. These papers deal with 'class' (defined in various ways), 'geography' (meaning location) and changes in inequality since 1990, to sustain arguments about changes in inequality during the contemporary period of globalisation and the distribution of the benefits of growth during that period. It is these papers that in part inspired the current paper's investigation. These papers are (in no particular order): Milanovic (2011; 2012), Palma (2011) and Ravallion and Chen (2012). Later we also note those papers focusing on the 'middle classes' such as Kharas (2010), Kharas and Gertz (2010) and Ravallion (2010) among others.

It is noteworthy that these papers, and the literature more generally, show a discernible shift away from reliance on single-figure measures of inequality, such as the Gini coefficient, and towards greater use of fractile (quintiles/deciles/ventiles etc.) data to explore how economic growth affects individuals at different consumption/income levels both within countries and transnationally.[4]

These recent studies noted above have made two arguments in particular: first, that 'global inequality' (defined in different ways) is falling because international inequality

(between countries) 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 are difficult to discern regionally with clarity.[5] This points towards the issue of the interplay of inequality between and within countries and the role of China in particular.[6] As we show below, rapid economic growth in China has been the dominant factor in recent reductions in population-weighted between-country inequality, and at the same time inequality within China has increased very significantly, influencing within-country inequality estimates. One question is thus if China is excluded then from 1990 to 2010, what has happened to inequality between and within countries across the rest of the world?

2.2 INEQUALITY, 'CLASS' AND GEOGRAPHY

The papers noted above approach inequality between and within countries in different ways. For example, Palma (2011) focuses predominantly on similarities and differences in within-country inequality across the world. Using a World Development Indicators dataset that includes observations for 135 countries (ibid: 89) with information on Gini and income shares, Palma reaches the following conclusions: first, he shows that about 80 per cent of the world's population now lives in regions whose median country has a Gini close to 40, implying that globalisation has reduced regional differences in within-country inequality. Second, Palma shows that the high inequality 'outliers' to this tendency are now only located among middle-income and rich countries. In other words, the 'poor and upwards' side of the Kuznets 'inverted-U' between inequality and income per capita has evaporated —and with it, Palma argues the hypothesis that posits that for poor developing countries inequality has to worsen in the earlier stages of development.[7]

Palma's third conclusion, and the one that has received attention, is the startling finding that within a global trend of rising inequality, 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 per cent (deciles 5 to 9). The other is 'centrifugal' and leads to an increased diversity in the shares of the top 10 per cent (decile 10) and bottom 40 per cent (deciles 1 to 4). The share of 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 per cent and most values within the range of 50–55 per cent (Palma, 2011: 101 and 102).[8] 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; the other half, however, is increasingly up for grabs between the very rich and the poor" (ibid: 103–104).

The remaining half of the national income is variously distributed between the poorest 40 per cent (deciles 1 to 4) and the richest 10 per cent (decile 10) in each country. There is, however, a wide diversity, between countries, in the share appropriated by the richest 10 per cent. 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 gross national income (GNI) of the richest and the poorest, and globalisation is thus creating a distributional scenario in which what really matters is the income share of the rich (because the rest 'follows').[9]

Palma points to three issues: first, the 'homogenous middle' outlined above. Second, that while most regions and countries have generally similar levels of inequality, two middle-income regions (Southern Africa and Latin America) have remarkably high levels of inequality, representing what probably amount to the most extreme practicable divisions between the rich, median and poor segments of society (since 60 is the maximum likely Gini value, one might speculate that while more extreme divisions are theoretically possible, they are likely to prove difficult to sustain consensually as functioning social systems, since they imply such a wide level of difference between the top and bottom). Third, it is among the richest countries that the highest diversity of distributions occurs (QED). High levels of development can (at least currently) co-exist with either high or low levels of inequality.

Milanovic (2011; 2012) in contrast, uses data (2011: 7) derived from the World Income Distribution database[10] 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. Milanovic calls this 'Inequality 3' and outlines it as follows:[11] "*Inequality 3* is [the] global inequality, which is the most important concept for those interested in the world as composed of individuals, not nations. Unlike the first two concepts, this one is individual-based: each person, regardless of her country, enters in the calculation with her actual income."

And Milanovic's (2012) central argument derived from this 'Inequality 3' concept is as follows:

> "We see something that may be historically important: perhaps for the first time since the Industrial Revolution, there may be a decline in global inequality. Between 2002 and 2008, global Gini decreased by 1.4 points… [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. The top 1% has seen its real income rise by more than 60% over those two decades. The largest increases however were registered around the median: 80% real increase at the median itself and some 70% around it. It is there, between the 50th and 60th percentile of the global income distribution that we find some 200 million Chinese, 90 million Indians, and about 30 million people each from Indonesia, Brazil and Egypt. These two groups —the global 1% and the middle classes of the emerging market economies —are indeed the main winners of globalization… So who lost between 1988 and 2008? Mostly people in Africa, some in Latin America and post-Communist countries" (2012: 7, 8, 12, 15).

Milanovic (2012: 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') in 1870 and 2000 and notes, "What is less obvious and less well known is that the shares of the two factors determining global inequality have changed in a remarkable fashion... 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" (2012: 19).

In short, Milanovic argues that global inequality has become much more about 'location' rather than 'class'. Milanovic's use of the term 'class' is, however, rather problematic. 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.[12, 13]

Although, as indicated above, we would suggest 'class' may not be a useful term to adopt when looking at global income distributions, 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.[14]

Finally, Ravallion and Chen (2012) provide a further analysis of inequality between and within countries by applying 'the mean-log deviation' (MLD) inequality measure to the PovcalNet database (of more than 850 household surveys from 127 developing countries) to explore changes in inequality within and between countries 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.: 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: 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. This is also evident if one takes out China, which has a high weight, and also saw a large increase in inequality… 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 short, 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 in recent decades —namely, that since the early 1980s within-country inequality has become an increasingly important element in overall global inequality.[15] It seems possible, therefore, that the long-term trend of the 20th century towards lower inequality within countries but higher inequality between countries could be starting to reverse, 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, caution is needed here because (as we show later) these global differences can be largely attributed to a few large emerging economies which disguise the fact that large differences between the remaining countries remain rather intractable (for example, when China is removed from the analysis we find that between-country inequality has risen since 1980, while within-country inequality hardly changed).

3 THE GRIP (GROWTH, INEQUALITY AND POVERTY) MODEL

In this paper we use a model of growth, inequality and poverty —the GrIP model, which has been developed from an earlier model described in Edward (2006). The GrIP model in its current iteration is described in detail in Edward and Sumner (2013).[16] The main objective of the GrIP model is to construct a truly global model of or consumption distribution that allows ready comparison of different assumptions. The model is built by combining consumption (or income) distribution data derived from national surveys with internationally comparable measures of mean national consumption.

The core approach in the GrIP model is to take for each country the distribution (quintile and decile) data and combine these with data on national population and on the mean consumption per capita in internationally comparable PPP $. The model uses this input to develop for each country an estimate of how many people live at any specific consumption ($/day) level. This is very similar to the approach in Povcal except that: GrIP uses a linear rather than kernel distribution function; draws on a wider range of sources to extend the analysis to include developed economies (including allowing estimation for countries where some of the source data are missing); and allows the use of a variety of possible sources for the mean per capita consumption (whereas Povcal relies only on survey means, which can limit the ability to extend the scope to include other countries). 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 available from a number of sources, although the most commonly used source is the World Bank's PovcalNet.[17] In GrIP, distribution data are taken (in this order of preference) from PovcalNet, World Development Indicators or the United Nations University World Income Inequality Database (UNU WIID – V2.0c, May 2008).[18] National consumption means can come from survey data or from National Account (NA) measures. The GrIP model enables us to make analyses using either survey means (Option 1 in GrIP) or NA means (Option 2). Survey means are taken from PovcalNet, while NA means are taken from World Development Indicators. All analysis and results are in 2005 PPP $.

There are a number of methodological issues to consider and acknowledge. First, even though these datasets have greatly improved their global coverage in recent years, there are still some significant gaps in the data so that, to construct a truly global distribution, it remains necessary to decide how to deal with missing 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. This still leaves situations where a country has no surveys or the gaps between surveys are too great to allow reliable interpolation. In these cases the GrIP model can estimate, or 'fill', a country's missing distributions with the (not population-weighted) average distribution from all other countries in the same region and income group (i.e. in the GrIP model the analysis can either be 'filled' to include these estimates and hence build a model that more closely replicates global population and consumption totals, or 'not filled', which means that the analysis only includes the smaller set of countries for which national distribution data are available and hence covers a smaller proportion of global totals).[19]

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

GrIP is designed so that the difference between using NA means or survey means can be readily investigated. Survey means are used, for example, in Povcal, but other authors tend to use NA means. One reason for this is that the NA means allow one to bring together distributions from a wider range of sources. By using the NA means as an internationally standard metric across which to aggregate and compare the individual countries, it becomes possible to incorporate countries where distribution data are available but comparable survey means are not available. It also means that informed estimates can be made for inequality in countries where distribution data are absent and so build a more truly global model. A difficulty arises because, although in theory survey means should show a consistent relationship to NA means, in practice this is not the case. A general systematic relationship (the NA/Survey ratio) can be estimated, but individual countries can show very wide variation from this estimate. Use of survey or NA means can lead, therefore, to a very different geography of global poverty and inequality, so it becomes relevant and useful to compare how the selection of the different means can affect results of the analysis.

When we use NA means, we simply apply the relevant mean to the survey distribution (this is termed NA/S Option 2 on the various graphs that follow). When we use survey means (Option 1), we apply the survey mean to the distribution where such a mean exists. Where there is no survey mean, we adjust the NA mean for the country in question in line with estimates for the systematic difference between the two types of mean (the NA/Survey ratio — NA means are in general higher than survey means). Because of this adjustment in Option 2 we also adjust by the same ratio the various thresholds between 'layers' that we identify later. For a more in-depth discussion of these issues see Edward and Sumner (2013).

In this paper we use Household Final Consumption (HFC) per capita means (in 2005 PPP $) as our NA mean data. Because coverage of GDP data is generally better than that of HFC data, where GDP data exist but HFC data do not, then the missing HFC figure is estimated from the GDP data. Wherever possible this is done in a given year by applying the most recent HFC/GDP ratio for the country in question. Where no such ratio exists, then the average ratio calculated for all countries with suitable data in the same region and income category is used.

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.

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. (%)	Consum-ption cov. (%)	No. of countries	Pop'n cov. (%)	Consum-ption cov. (%)	No. of countries	Pop'n cov. (%)	Consum-ption 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: GrIP v1.0. Percentages are of global totals.

4 WHAT HAS HAPPENED TO INEQUALITY BETWEEN AND WITHIN COUNTRIES SINCE 1990?

4.1 GLOBAL INEQUALITY AND DECOMPOSITIONS

Two commonly used measures for global inequality are the Gini index and the Theil index. Of these, Gini is the more widely used, largely because of its close and relatively intuitive association with the Lorenz curve. However, Gini suffers from the problem that it is not decomposable, so it becomes difficult to differentiate the separate contributions of within-country and between-country inequality to overall global inequality. In essence, in a truly global model of inequality, changes in the global Gini coefficient arise from three causes: a) changes in between-country inequality (Milanovic's population-weighted 'Inequality 2'); b) changes in within-country inequality; and c) the interaction of (a) and (b).[20]

It is, therefore, difficult to know whether to ascribe the interaction, or overlap, element (c) to between-country or within-country changes. But this element can be significant, especially in situations where highly populated countries are experiencing rapid changes both in aggregate national consumption and in within-country inequality (as has been the case recently with the large emerging BRIC economies —Brazil, Russia, India and China). 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.

Gini and Theil both give us measures of global inequality, but their different sensitivities mean that they are not directly comparable. For this reason we provide estimates here using both indices for the period from 1980 to 2010 (noting again caveats for the 1980s). We also include for comparison estimates from Milanovic (2012b: 14) based on survey means. Many other earlier estimates of Gini and Theil exist, but they are not directly comparable to our analysis because, as Milanovic shows, the recent rebasing of international PPP rates as a result of the 'new' 2005 International Comparison Program (ICP) revision (relative to the 'old' 1993-based PPP rates) has systematically increased global inequality indices.

Gini and Theil results are shown in Figures 1 to 6 (all these coefficients are based on 'not filled' analysis so that they are not distorted by any estimations for 'missing' countries in the filling process). We also include for the Theil index graphs showing how within-country inequality has changed as a proportion of total global inequality. For all these figures we show results based on survey means (Option 1) and on NA means (Option 2).

We can see from the Gini figures that global inequality has certainly declined since the late 1980s, but the amounts are relatively small. Perhaps it is more accurate to say that global inequality by the Gini method was relatively constant through the 1990s (the variability in the figures prior to the early 2000s looks to be random scatter within the margin of error, which would also encompass Milanovic's estimates, rather than an overall trend). Since the early 2000s, however, global inequality has been falling. The causes of this are interesting, because across the entire period within-country inequality has been rising slowly but steadily. In the 1990s this rise was approximately balanced out by a gradual fall in between-country inequality. Since the early 2000s the rise in within-country inequality has continued at a similar steady rate, but 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 started to fall.

The Theil index gives a broadly similar conclusion. Its greater sensitivity to changes at the extremes of the distribution means that the results based on survey means might be read as indicating that global inequality was rising slowly in the 1990s rather than broadly static, but that effect is not robust enough to show up also when NA means are used.

In sum, taken as a whole the results indicate that for the first 10 to 15 years of the period of globalisation since the late 1980s, global inequality was relatively static, with a slow but steady rise in within-country inequality being broadly offset by a gradual decline in between-country inequality. Since then, and particularly since 2005, while within-country inequality has continued to rise steadily, 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 20–25 per cent of global inequality, by 2010 it had risen to 30 per cent of global inequality.

The figures are consistent with Ravallion and Chen (2012) and seem to be the reverse of the longer-term trend Milanovic identifies since 1870 (2012: 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.[21] While one should be cautious about relying too heavily on the data from the 1980s, interestingly it indicates that the rise in within-country inequality started in the late 1980s (around the time of the fall of the Berlin Wall and the ascendancy of market liberalism) and has continued steadily since then. Between-country inequality, however, was fairly static until the early 2000s when it started to fall quite quickly. As a result, overall global inequality seems to have risen from the early 1980s to the early 2000s but has been falling quite sharply since then.

However, the picture looks rather different when China is excluded (Figures 7 to 9). 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. The rise is particularly noticeable in the Theil coefficient, which might indicate that a lot of the rise was caused by increasing inequality between the richest and poorest countries (since Theil is more sensitive than Gini to changes at the ends of the distribution). Since 2000 between-country inequality has been fairly static (when China is excluded), but there is little sign that a reversal has set in following the developed world's financial crisis in 2008. Perhaps it is too early anyway to detect the effects of that crisis in the data, but 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 since 1980 of increasing inequality between rich 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 should give us pause for thought before celebrating too keenly recent and very modest signs of falling overall global inequality. The rapid progress of China may be masking underlying trends that are rather less progressive.

FIGURE 1

Global Gini Coefficient, Survey Means (not filled)

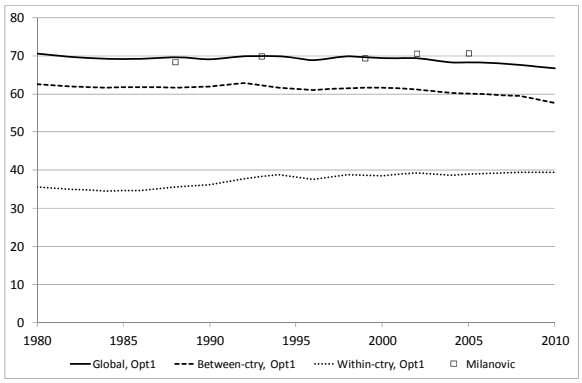

FIGURE 2

Global Gini Coefficient, NA Means (not filled)

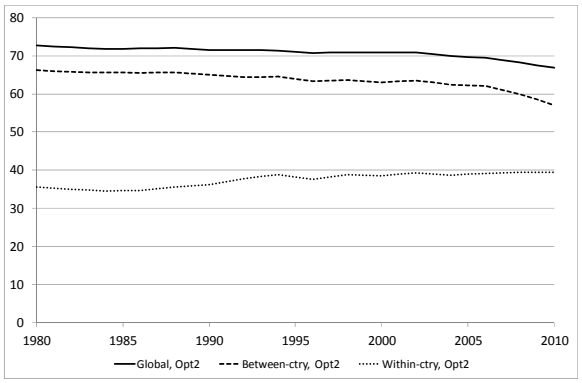

FIGURE 3

Global Theil Index, Survey Means (not filled)

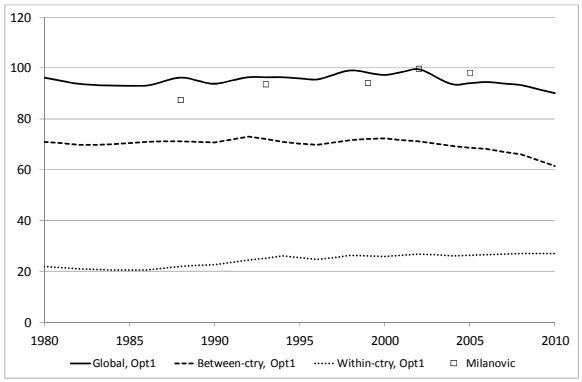

FIGURE 4

Global Theil Index, NA Means (not filled)

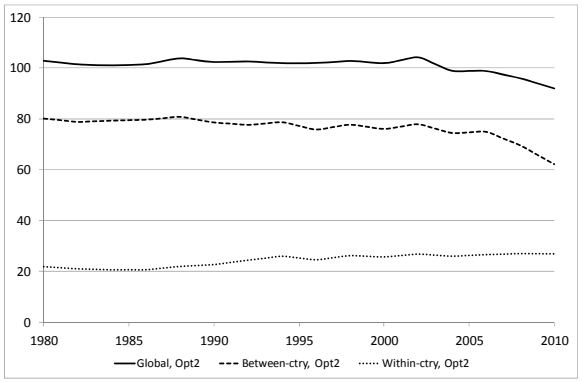

FIGURE 5

Within-country Theil Component as Percentage of Global Theil Index, Survey Means (not filled)

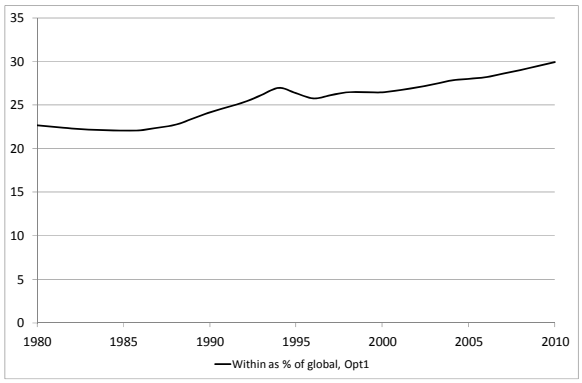

FIGURE 6

Within-country Theil Component as Percentage of Global Theil Index, NA Means (not filled)

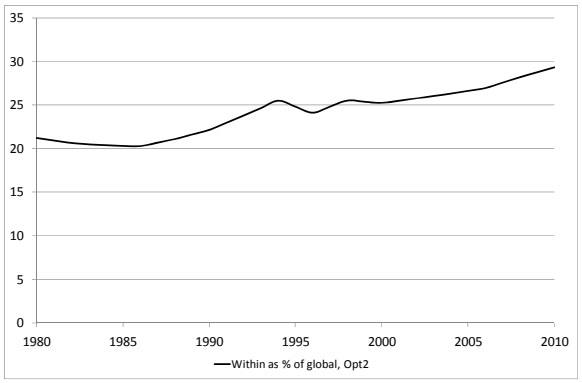

FIGURE 7

World excl. China, Gini Coefficient, Survey Means
(not filled)

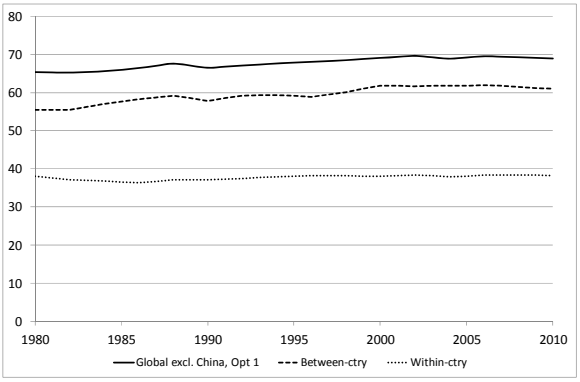

FIGURE 8

World excl. China, Theil Index, Survey Means
(not filled)

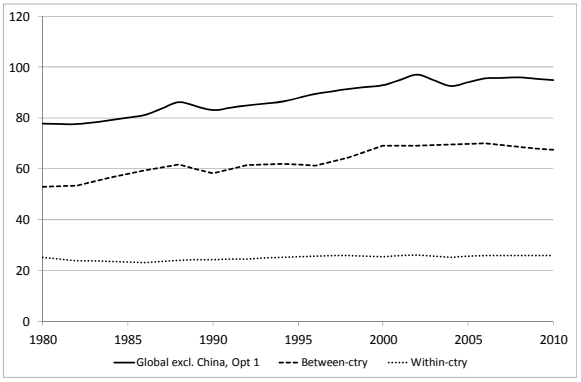

FIGURE 9

World excl. China, Within-country Theil Component as Percentage of Global Theil Index, Survey Means
(not filled)

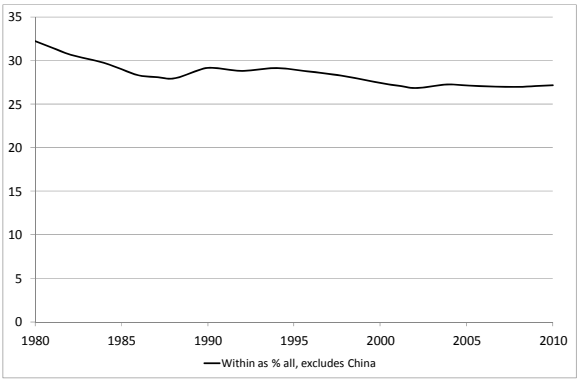

4.2 WITHIN-COUNTRY GINIS BY REGIONAL AND COUNTRY GROUPS

One way to look at within-country inequality further is to look at how it has changed. What do the data say? Figures 10 to 14 show, respectively, Gini coefficients by regional and country groups and selected individual countries. All are survey means (Option 1) and 'not-filled'. Regional groupings do not include high-income countries (HICs) within that region.[22]

As Palma (2011) found, Latin America and sub-Saharan Africa are the most unequal regions in the world. As others have noticed, however, in Latin America since 2000 inequality has been falling, reversing a trend of rising inequality in the 1990s. This is largely the result of falling within-country inequality in the region. While Brazil has been a significant contributor to this fall, it has not been confined to Brazil, since inequality has also been falling across the rest of the region since 2000 (See Annex Table A1). However, it remains to be seen whether this trend will continue or whether it represents merely the latest phase in the cyclical fluctuation of the regional Gini around a value of 50 to 55. At first sight, sub-Saharan Africa's fluctuations seem to be cyclical also. However, there are some significant differences. Within-country inequality has fallen across the region as a whole since 1990 despite a significant and

continuing increase in South Africa. When South Africa is removed from the analysis, the picture for the rest of sub-Saharan Africa is one of relatively constant overall inequality, as rising between-country inequality is offset by falling within-country inequality.

In East Asia overall inequality has risen sharply since the late 1980s, driven not surprisingly by a rise in within-country inequality that is largely down to changes in China. When China is removed from the analysis, within-country inequality in the region is found to have risen only slightly since the mid-1980s. Furthermore, to put some context on this, China's inequality appears even now to be both slightly lower and growing less rapidly than the USA's. By all measures, South Asia remains one of the least unequal regions (which might not be a good thing if this merely reflects high levels of absolute poverty), but there are signs that inequality there may be starting to increase.

FIGURE 10

Gini Coefficients (between- plus within-country inequality) **by Region**

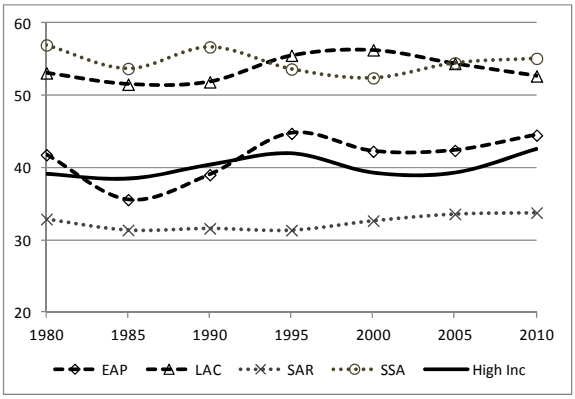

FIGURE 11

Gini-coefficient for Sub-Saharan Africa, Excluding South Africa

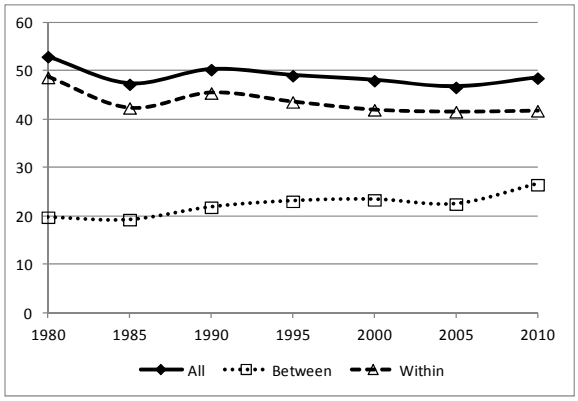

FIGURE 12

Gini Coefficients (within-country inequality only) **by Region**

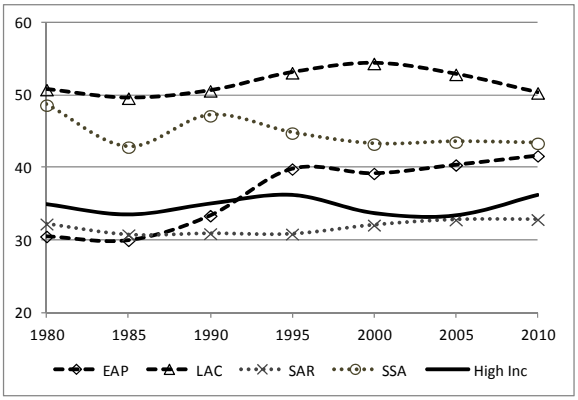

FIGURE 13

Gini-coefficient for USA and EU

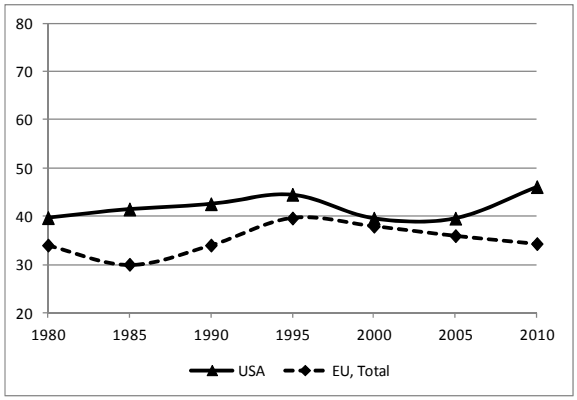

FIGURE 14

Gini Coefficients for Selected Countries

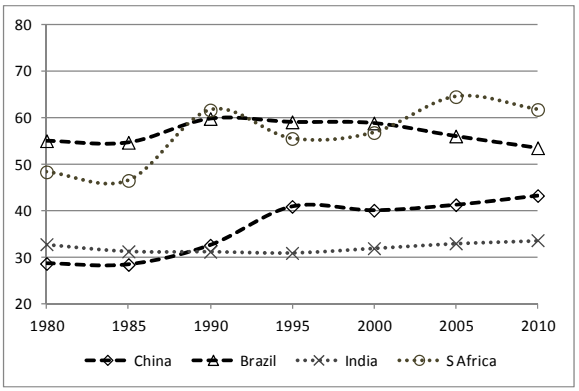

5 INEQUALITY IN A MULTI-LAYERED WORLD

5.1 THE 'LAYERS' OF GLOBAL SOCIETY

In this section we consider a different angle, focusing on inequality in a multi-layered world. In short, we consider the fortunes of 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'.

One conclusion to draw from the inequality indices 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. Who have been the winners and losers, for example, from the period of global economic expansion that has taken place since the fall of the Berlin Wall in the late 1980s? These issues become rather lost when we focus on simple indices and on individual countries.

A true global inequality picture (such as Milanovic's 'Inequality 3') can, however, allow us to move beyond these limitations and to develop a properly transnational view of changing global inequality. One approach has been used for several decades now to estimate global $/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). The potential of models such as GrIP is that if they incorporate survey data that cover all countries (and not just the poorer countries) then one can start to look at consumption levels and categories as a global phenomenon.

Ravallion starts to do this, although only for developing countries, by defining the middle class as the population segment between a lower-bound absolute threshold of $2/day and an upper bound threshold of $13/day (the US poverty line) (all in 2005 PPP $). By contrast, Palma and Milanovic tend to understand class as a relative, national issue. Palma defines the 'middle class' as those between the 40th and 90th percentiles. Milanovic does not define classes but writes of within-country inequality as derived from class, whereas between-country inequality is derived from location. Elsewhere Kharas, for example, has defined the middle class against absolute thresholds, with a lower bound of $10/day and an upper bound of $100/day (Kharas, 2010a; 2012b).

This paper uses absolute thresholds to identify a multi-layered global society, effectively taking an approach which has an established history from the use of global poverty lines and extending it to cover the full range of global consumption levels.

However, as noted above, we dislike the use of the term 'class' here and instead propose to call these divisions consumption 'layers'. Our reasons are that class itself is a social and political identity not necessarily linked to estimates of expenditures per capita. Although 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 this 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). We, therefore, find the notion of transnational classes to be unhelpful as a basis for trying to identify alignments in consumption levels between different segments of global society

Segmentation by absolute consumption levels, as we use here, is limited to the extent that we are grouping people globally by their consumption levels —the kind of lifestyles they lead —rather than by any deeper 'class'-derived alignment of socio-political 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).

Class is, of course, most often discussed in terms of types of assets and productive processes, labour markets and occupational resources (see review in Torche and López-Calva, 2011). Some contemporary sociological analysis of class also places a particular emphasis on economic security (see, for discussion, Erikson and Goldthorpe, 2008; Goldthorpe and McKnight, 2006; Standing, 2011). When we extend these largely economic categorisations transnationally, the basis for any alignment becomes all the more one of similar material, rather than cultural or political, interests. Collectively, therefore, these studies provide justification for using absolute levels of consumption to segment global society on the basis of economic security (vulnerability to poverty) at lower levels of consumption and of broad similarity of consumption preferences at higher levels of society —but they provide little justification for describing these segments as global 'classes' —that is, as united by a shared socio-political identity.[23]

In this paper we propose, instead of class, four consumption levels or 'layers' 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 clusterings of consumption levels, 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 (ie. similar spending or consumption power), then we might think of that as a distinct global cluster. If such a group existed, then it would be clustered

(presumably with some sort of vaguely 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 per cent of the distribution would lie. This 85/15 division makes a useful criterion for selecting thresholds.

Consumption density curves from the GrIP model are this sort of plot. In Figure 15 we present the density curve for the world in 2010. In the graph consumption per capita (2005 $ PPP) is plotted on a log scale on the horizontal axis. The vertical dashed line (at $730 per annum) represents the $2/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/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 (for the statisticians it is normalised so that the entire area bounded by the population curve and the horizontal axis aggregates to unity).[24]

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). So, when we plot other population curves in this paper, their areas are all relative to and in proportion to the 2010 global population curve. Similarly, consumption curves are all relative to the 2010 global consumption curve (Option 1 and Option 2 graphs are relative to their respective 2010 curves).

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 in Figures 16–19. On these plots we show our three proposed thresholds and also the World Bank's current extreme poverty line ($1.25/day). These plots are included mainly to show that the basis for our thresholds is broadly robust to changes in modelling assumptions (use of NA or survey means). Here, the 'filled' analysis is utilised so that the curves represent similar proportions of global population. In this case they cover 96–97.5 per cent of the entire global population (see earlier discussion and Table 1).

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, as described above. Our segmentation approach approximately divides each of these clusters into an upper and lower segment, and identifies a dividing point between the two clusters.

FIGURE 15

Global Density Curve, Survey Means, Filled

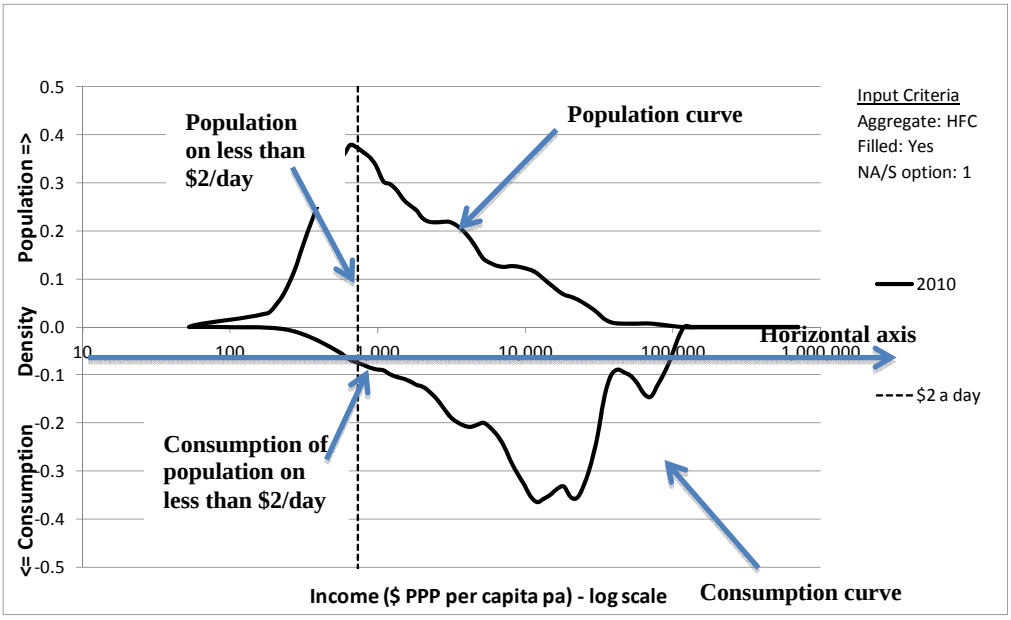

FIGURE 16

Global Density Curve, Survey Means, Filled

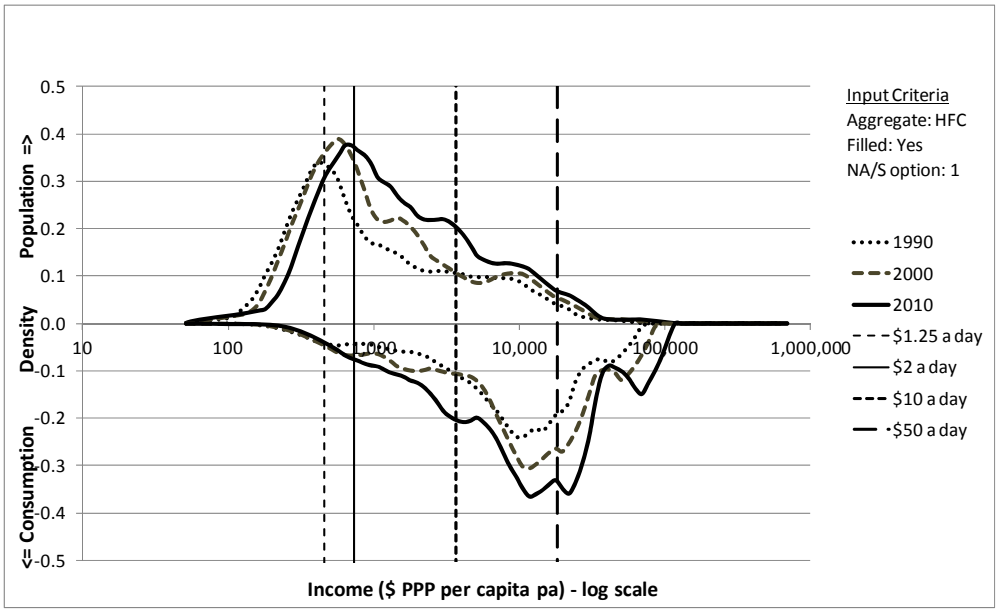

FIGURE 17

Global Density Curve, NA Means, Filled

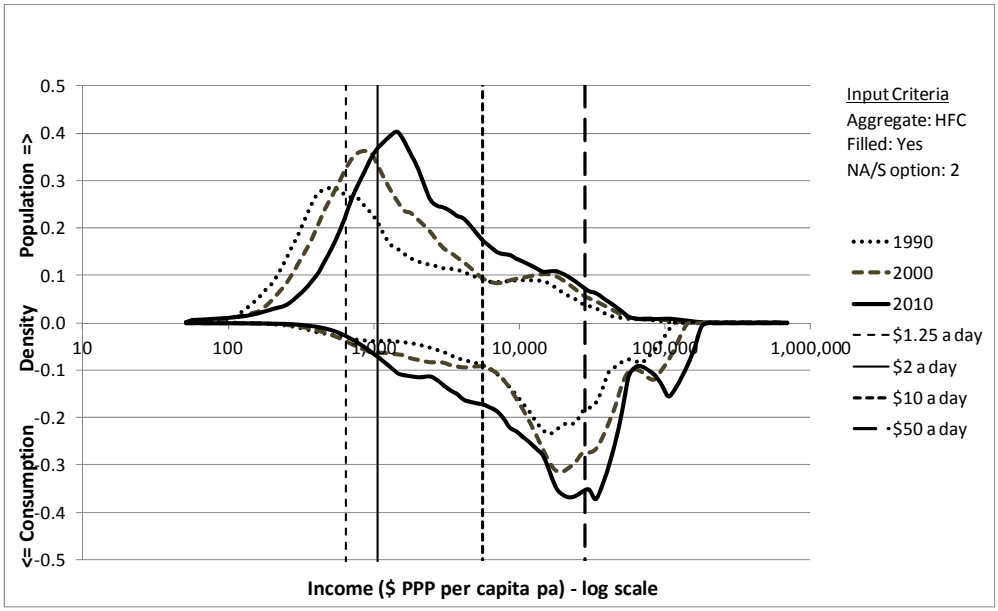

FIGURE 18

Global Density Curve, Survey Means, Filled – Excluding China

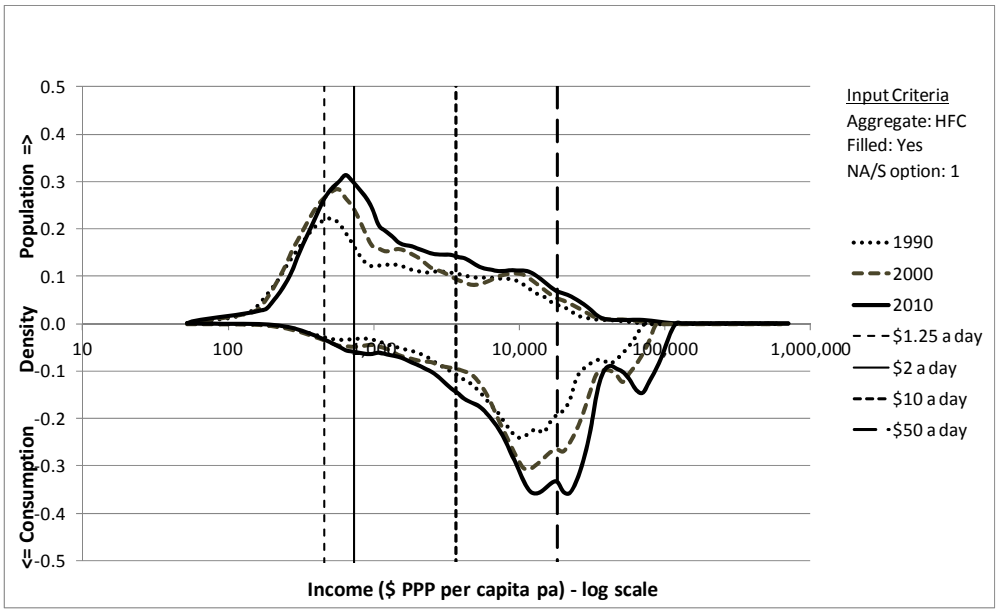

FIGURE 19

Global Density Curve, NA Means, Filled – Excluding China

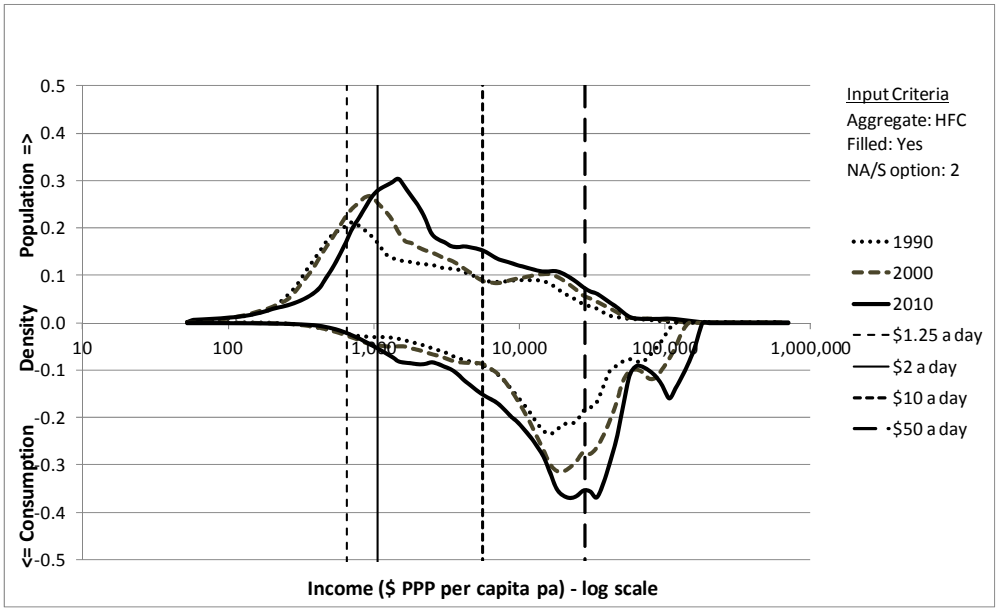

We, therefore, derive our four global consumption 'layers', or segments, as follows: first, the 'global absolute poor' —we define this as living under the $2 per capita level. This is not only reasonably close to the mid-point of the low-income peak, it is also well established as the World Bank moderate international poverty line, which is close to the median poverty line across all developing countries ($2.36 per capita in 2008) as well as the regional mean poverty line in sub-Saharan Africa and the South Asia Region and China, collectively where many of the world's poor people live (Ravallion, 2012: 25). The global mean for poverty lines for developing countries is $4.64 per capita, which is rather higher than the median, because poverty lines can be around $11–12 per capita —the mean in Latin America and the Caribbean and in Eastern Europe —and close to $4 —the mean in East Asia and Pacific (ibid.). In short, $2 per capita seems reasonable because it is close both to the global median and to the mean poverty line in the countries where most of the world's poor people live (sub-Saharan Africa, South Asia Region and China). It is certainly more appropriate than the extreme $1.25/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/day represents a reasonable cut-off point in the overlap between these two peaks. From GrIP, 87 per cent of the population of HICs lives above $10/day per capita —which fits our 85/15 rule —while 98 per cent of populations of low-income countries (LICs) and lower middle-income countries (LMICs) 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 per cent of the HIC population above the threshold and 94 per cent of LIC and LMIC populations below it. Although any line is arbitrary to some extent, there are reasons not to go any lower than $10, since $10 has been identified as an approximate level of 'security' from poverty.

An interesting study in Chile, Mexico and Brazil suggest that the risk of falling into poverty was as low as approximately 10 per cent at an initial income of $10/day per capita in all three countries but fell to zero in Chile and Mexico at an initial income close to $20/day. The authors refer to this as a "vulnerability approach to identifying the middle classes" (López-Calva and Ortiz-Juarez, 2011). And 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. Ravallion (2010) used an even higher threshold, the US poverty line of $13/day per capita. 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 above $50/day per capita. This approximates to the mid-point of the higher-income peak (so that around half of HIC consumption is by people living above this level, and half is below it). It is also the level below which 87 per cent of the HIC population lives (so fits our 85/15 'rule'). A further reason for choosing this location as the division between the global secure and the global prosperous is that it neatly 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.

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/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. Figures 20–22 show the distribution of global society by the layers outlined (see also Annex Tables A2 and A3).

FIGURE 20

Estimates of each Layer of Global Population (millions)
by Region, 1990 and 2010: HFC, Survey Means, Filled

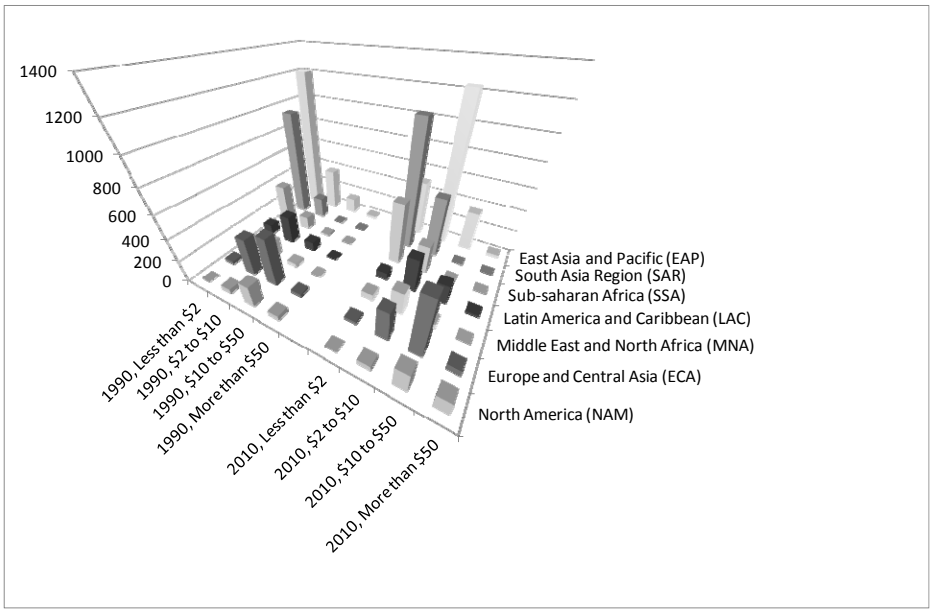

FIGURE 21

Estimates of each Layer of Global Population (millions)
by Country Income Groups, 1990 and 2010: HFC, Survey Means, Filled

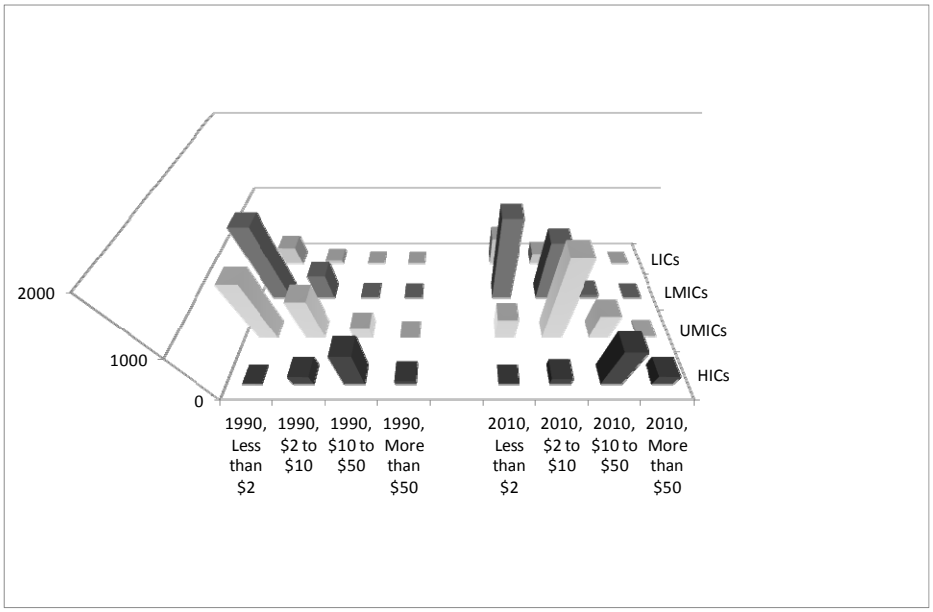

FIGURE 22

Estimates of each Layer of Global Population (millions)
for Selected Countries, 1990 and 2010: HFC, Survey Means, Filled

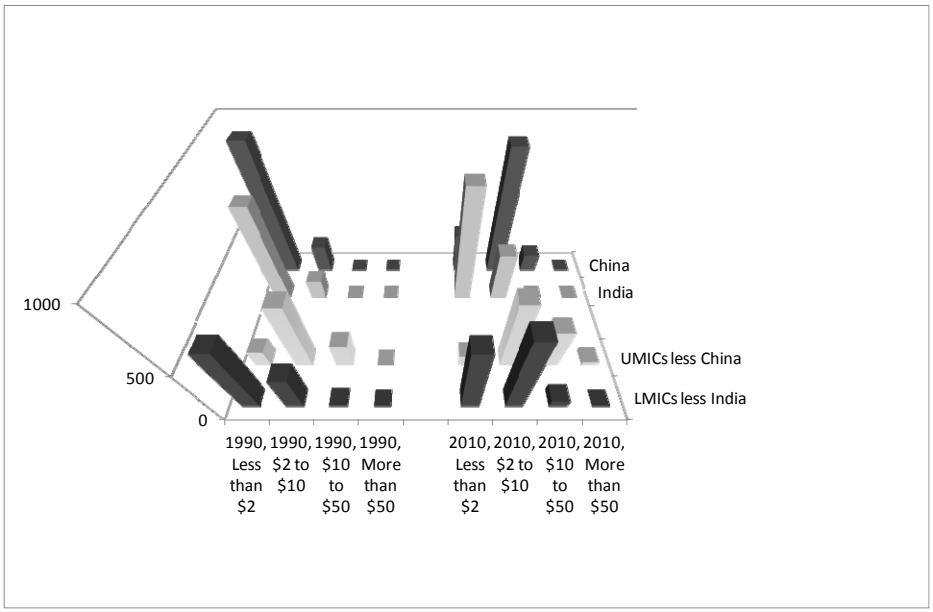

5.2 TRENDS IN THE CONSUMPTION DISTRIBUTION

The 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 23). 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 'dumbbell' 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 dumbbell has become much less distinct (see Figure 24), best seen as the loss of concavity between the poor and rich peaks in the density curves, presented above and below, so that the dumbbell looks, in 2010, more like a rotated rectangle. [25]

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 this may not be what is happening, since once China is removed, the population curve remains concave at consumption over $2/day.

Another way of interpreting this is that in the twin-peak 'dumbbell' world of the 1980s (and, doubtless, previous to that) we lived in a world with 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.

FIGURE 23

The 'twin-peak' Global Distribution of the Mid-1980s

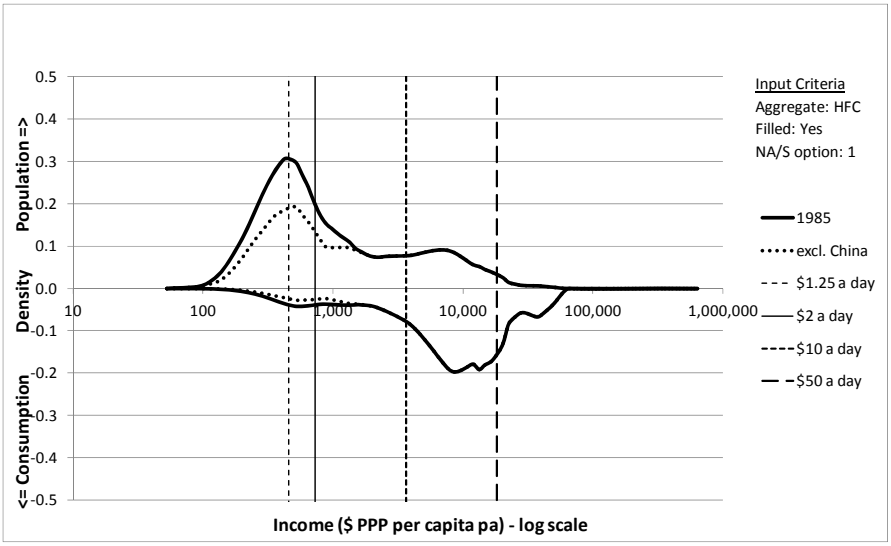

FIGURE 24

The 2010 Global Distribution and Loss of the 'twin-peaks' or 'dumbbell'

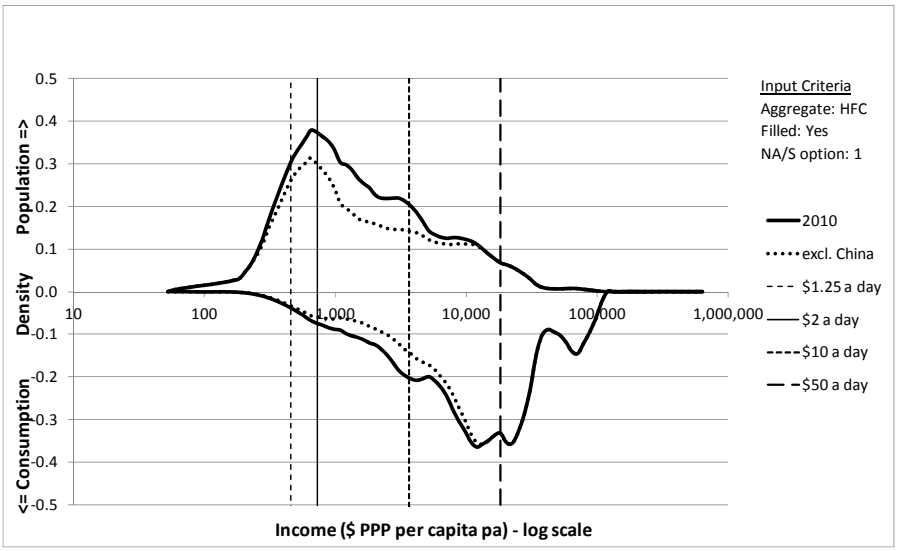

That the old 'dumbbell' structure largely persists is indicated in the 2010 curves, which still clearly show a distinct peak for the global poor (best seen on the population curve, albeit now at an income near the $2/day line, whereas in 1990 it was close to $1.25/day) and another distinct peak for the global prosperous (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 per annum axis were not plotted on a log-scale). In other words, rather than seeing the formation of a more globally integrated socio-economic 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/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 and previously in the poor/insecure cluster move across to the secure/prosperous cluster.

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 25–31 show density curves for current (2010) categories of HICs, UMICs, and 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 populations 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. Noteworthy also is that, despite overall economic growth and rising mean incomes, total numbers of people living in extreme poverty ($1.25/day) across the LICs and LMICs have not changed much since 1990.

FIGURE 25

Density Curves for Current HICs

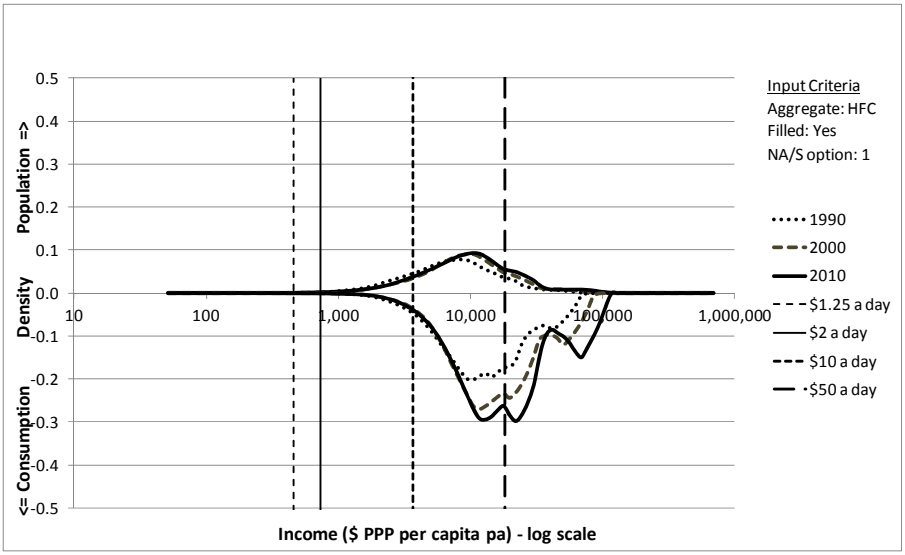

FIGURE 26

Density Curves for Current LICs and LMICs

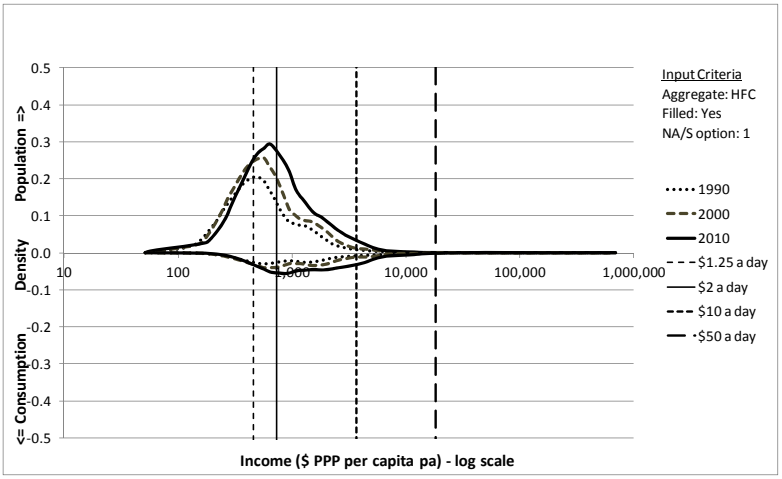

FIGURE 27

Density Curves for Current UMICs

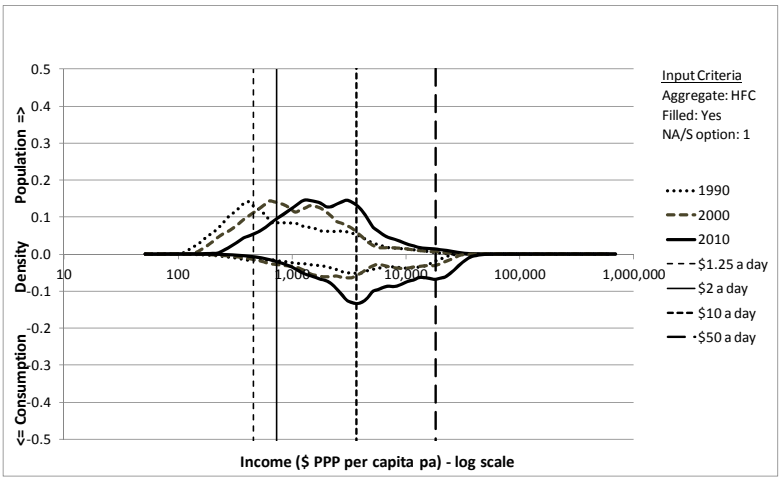

FIGURE 28

Density Curves for Current LICs and LMICs Excluding India

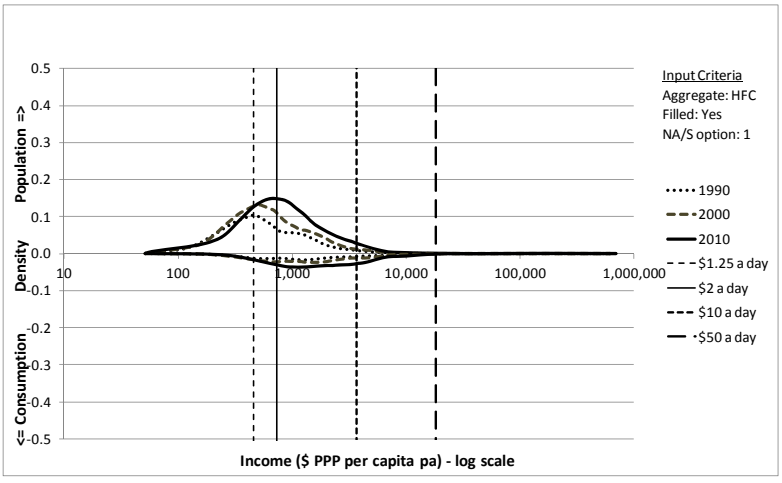

FIGURE 29

Density Curves for Current UMICs Excluding China

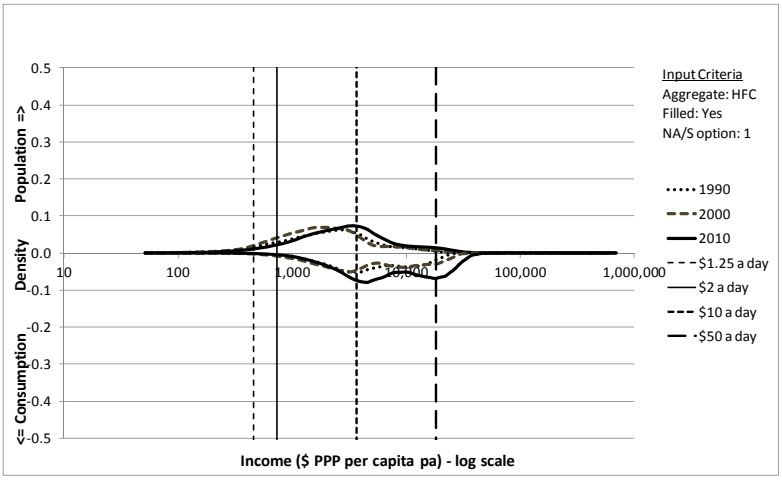

FIGURE 30

Density Curves for India

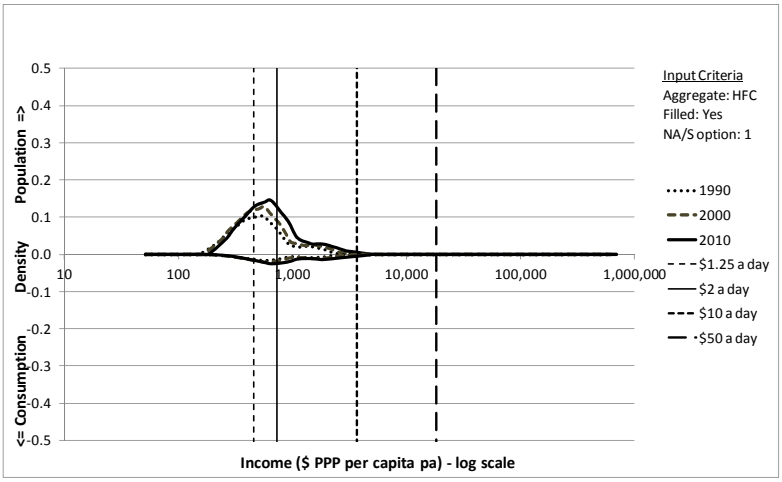

FIGURE 31

Density Curves for China

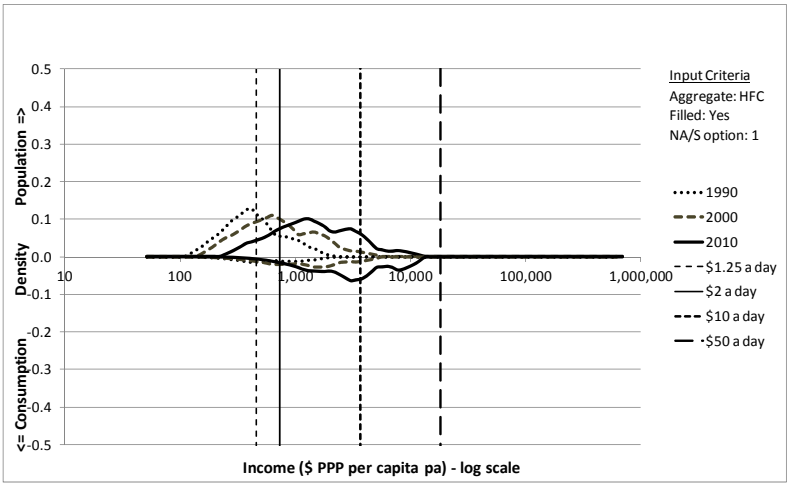

If we look at the HICs, there is little evidence of radical changes in the lower 'global secure' half of the distribution structure here (that is, in the half that would be affected if we were witnessing global convergence away from the twin-peak structure). The HICs still largely retain the presence of a relatively homogeneous 'global secure' component, with the main changes in these countries being some evidence of the evolution of greater differentiation among the upper 'global prosperous' levels in these countries (as discussed previously).

It is really only in the picture for the UMICs 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 'dumbbell' of a bipolar world of highly differentiated segments (the 'global absolute poor' and 'global insecure', on the one hand, and the 'global secure' and 'global prosperous', on the other hand) is seen to persist (Figure 32). 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 32

Density Curves for all Countries Excluding Current UMICs

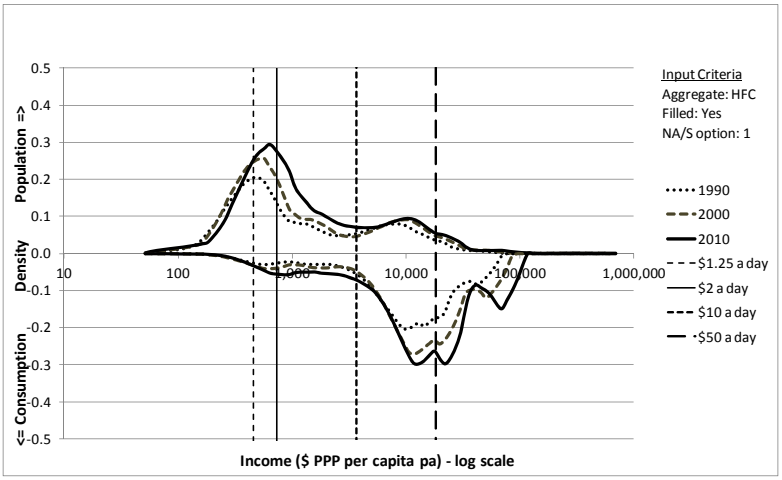

Similarly, in the aggregated LMICs and LICs there is still a reasonably well-defined peak. There is some evidence that wealthier elements may have been pulling away in the 1990s (the incipient formation of a second population peak in 1995), but the 2010 curve seems to be returning to a more homogeneous 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 HICs (and clustered around a median consumption of about $25/day per capita in 2005 PPP international $) and a relatively

poor grouping living in the LMICs/LICs (and largely clustered around a median consumption of about $2/day per capita in 2005 PPP international $ —or $1.70, to be more precise).

Certainly within China at present 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/day 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/day, a smaller middle cluster centred on $6.50/day and a richer cluster centred close to $18/day. The data suggest that 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).

5.3 WINNERS AND LOSERS FROM CONTEMPORARY GLOBALISATION

A further way to approach the impact of growth on global inequality is to consider by how much people living at different levels in the global consumption distribution have benefited or not since 1990. In other words, we can look at how global growth has been distributed among the various levels of prosperity. To do this, we present various charts that show how the average consumption at different consumption levels has changed from 1990 to 2010, a period that represents the ascendance of neoliberalism since the end of the Cold War. 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 the 50th percentile in 2010 has increased relative to that of the 50th percentile in 1990.

We present the data as charts where the horizontal axis represents fractile rank ordered by level of consumption. The vertical axis represents the percentage change in consumption per capita from 1990 to 2010. 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. Therefore, to provide a common reference point, we plot on all the graphs the average global per capita change in HFC, which grew by 47 per cent from $3930 per capita per annum in 1990 to $5770 per capita per annum in 2010. A great deal of information can be read from these charts, as Figure 33 demonstrates.

This analysis is, not surprisingly, quite sensitive to the model assumptions. In the density curves presented earlier it was important to use filled data so that the areas of each curve matched, as closely as possible, the global totals for population and consumption (for example, if we had used unfilled data, the 2010 consumption curve would have included only 80 per cent of global consumption, whereas the 1990 one would have included 93 per cent —rendering any comparison of the curves difficult). For the 'percentage change' graphs (Figures 34–37) the comparison between years is built into the analysis that generates the curves; therefore, we can use unfilled analysis to build the curves. In other words, these graphs do not rely on any estimates or 'filling' for countries that do not have available survey data, nor do they include a country unless there are usable data for both 1990 and 2010 (so the sample composition is identical in 1990 and 2010). In theory this is a dataset that is closer to the source data.

FIGURE 33

Illustration of Information in Fractile charts. Change in Consumption 1990 to 2010 at various Fractiles, Survey Means (Option 1), **not Filled – World**

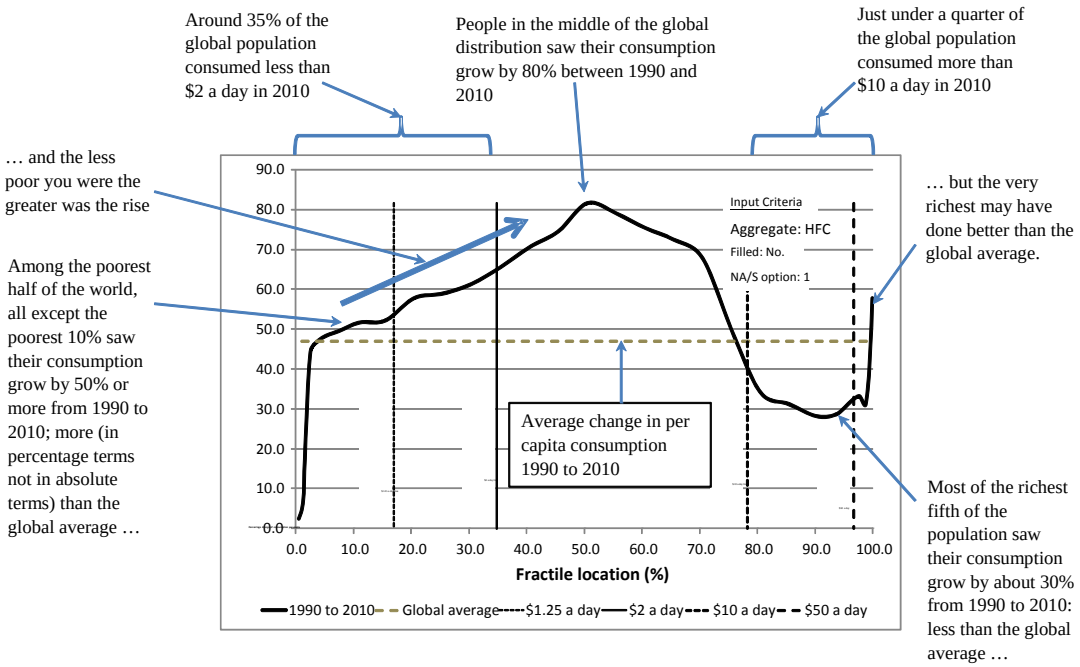

Nevertheless, the graphs do still need to be treated as indicative only because they may be distorted by the (lack of) availability of data. For example, in sub-Saharan Africa in 1990 GrIP includes surveys for 78 per cent of the population (and 93 per cent in 2010). So the percentage change graphs really refer to only about 80 per cent of the population of sub-Saharan Africa. Consequently, any systematic bias from the exclusion of countries without data, such as that they may be among the poorer countries in a region/aggregation, would influence the generalisability of these graphs to the whole region/aggregation. Among the regions included in this paper, the issue is most significant for sub-Saharan Africa, since in all other regions GrIP includes data for around 90 per cent or more of the regional population in both 1990 and 2010.[26]

On these graphs, distribution-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.

The use of survey means (Option 1) or NA means (Option 2) does affect the distribution of growth, with survey means estimating less benefit to the global poor than do NA means. This is consistent with findings that the use of NA means leads to estimates of faster poverty reduction and lower poverty headcounts than does the use of survey means (for a fuller discussion of this issue and the importance of recognising this systematic bias between different approaches to poverty and inequality modelling see Edward and Sumner, 2013).

Based on the survey mean charts we can see that the main beneficiaries of global growth since 1990 have been those in the 'global insecure' segment ($2–10/day) and below the 70th percentile. They saw their consumption rise by 60 per cent or more, considerably above the

global average. The main losers were the global secure ($10–50/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 per cent globally benefiting by more than the global average. Half of the 35 per cent of the global population that comprise the global absolute poor (below $2/day) are extremely poor (less than $1.25/day).

Nearly all of the absolute poor benefited at least in proportion to average global growth, but in general for the poorest 50 per cent 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.[27]

FIGURE 34

Change in Consumption 1990 to 2010 at various Fractiles, Survey Means (Option 1), **not Filled – World**

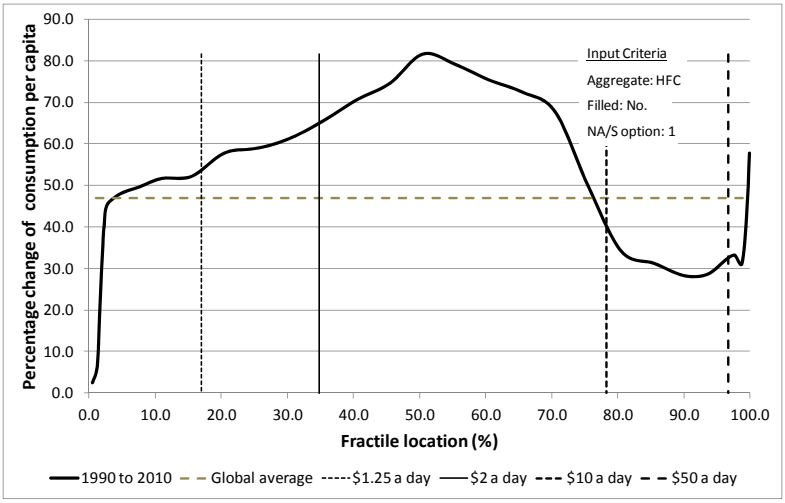

FIGURE 36

Change in Consumption 1990 to 2010 at various Fractiles, Survey Means (Option 1), **not Filled – World Excluding China**

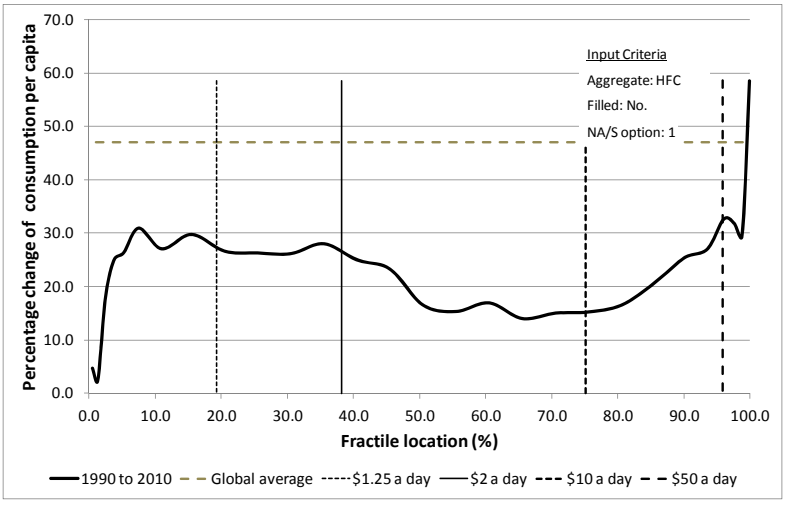

FIGURE 35

Change in Consumption 1990 to 2010 at various Fractiles, NA Means (Option 2)**, not Filled – World**

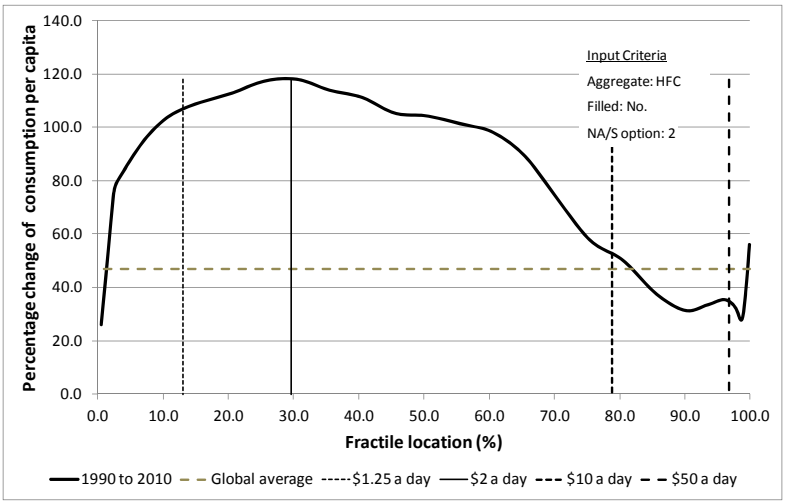

FIGURE 37

Change in Consumption 1990 to 2010 at various Fractiles, NA Means (Option 2)**, not Filled – World Excluding China**

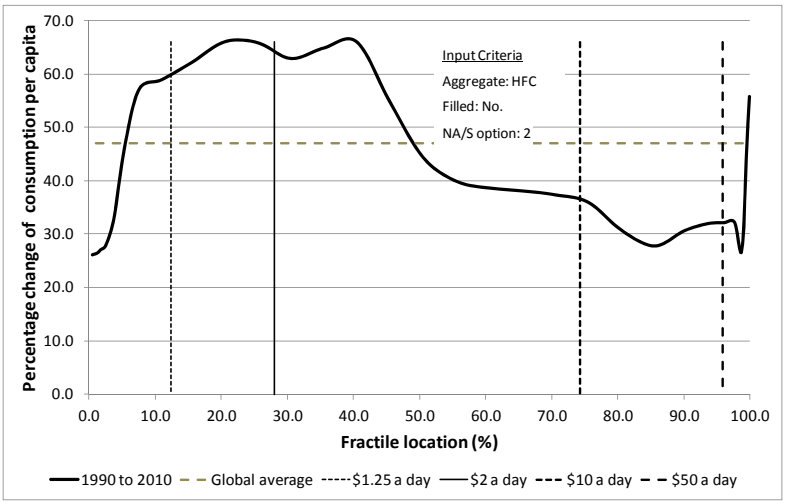

FIGURE 38

Absolute Change in Income Per Capita by Percentile – World

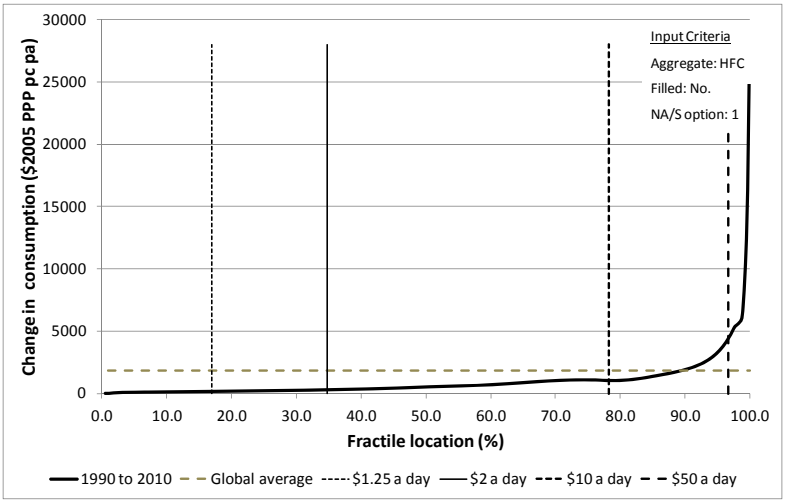

The picture, however, looks rather less progressive if China is excluded from the analysis (Figures 36 and 37). Most of the global absolute poor (outside China) saw their consumption levels rise by 25–30 per cent and so generally benefited more than the global secure and insecure, most of whom saw rises of around 15 per cent. However, the layer that benefited most were the global prosperous, fewer than 5 per cent of global population, who all saw their consumption grow in percentage terms by 30 per cent or more (and possibly for the richest by twice that amount), as much or more than any other segment of global society.

These figures can, however, be a little misleading about who really benefits from global growth, because they represent relative increases in consumption. An alternative way to look at this is to ask how individuals benefit absolutely —after all, 20 per cent of $25/day is a lot more than 60 per cent of $2/day. Figure 38 illustrates how in absolute terms only those in the world's richest 10 per cent have seen their consumption rise by more than the global average, with the main beneficiaries being the global prosperous and particularly those in the global top 1 per cent, on family incomes in excess of around $250,000 per annum —although caution is needed before drawing conclusions about this, the very richest percentile, because of the known inaccuracies of surveys noted earlier.[28] Or, to put it another way, global consumption grew by $10 trillion from 1990 to 2010,[29] and 15 per cent of that global growth accrued to the richest 1 per cent of the global population (see Tables 2 and 3), while the 3 per cent 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 per cent of global consumption growth. By contrast, the one third of the global population that are the now global absolute poor received just 5 per cent of the global growth, while the two fifths that are the global insecure received 25 per cent of global growth. This is *ex post* and based on the fact that there were 35 per cent of the global population on $2/day or less in 2010. The 5 per cent is the share of the total global consumption growth from 1990 to 2010 that went to the bottom 35 per cent (i.e. it is the difference between the aggregate consumption of the bottom 35 per cent in 1990 and the bottom 35 per cent in 2010).

An alternative perspective is to consider the *ex ante* case. Of the people living on under $2/day in 1990, what share of the $10 trillion did they get? In 1990 about half the world population was living on less than $2/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 living on less than $1.25/day —and collectively they benefited by little more than a 20th of global growth since then.

We might also ask how hard it really is to remove poverty. If we look at $2 poverty in 2010, the total poverty gap (estimated using survey means) is 6.6 per cent of the global growth from 1990 to 2010. In other words, between 1990 and 2010 the one third of global population who are living on $2/day today received 5 per cent of global growth, while 95 per cent went to the rest of the world. If in those two decades the world had been willing and able to pursue poverty-reducing policies strong enough to remove $2 poverty by 2010, then we would have needed to increase the share of growth flowing to the $2 poor by just 7 per cent of global growth. This would have meant that the non-poor would have received 88 per cent rather than 95 per cent of the global growth in that period. While this is a reduction, it is not very dramatic —almost certainly not so dramatic in itself as to call into question the logic of competitive self-interest that underpins capitalism. In short, it seems that in the last two decades there has been more than enough growth to remove poverty without substantially challenging modern aspirations to ever-increasing prosperity. The persistence of global poverty seems to have little to do with there being insufficient global growth and a lot more to do with a lack of collective will to forego some of the benefits of global growth in favour of modest redistribution.

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

Note: Numbers are derived using filled analysis so as to most closely match global aggregates; see also text and footnote 29 on top 1 per cent.

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

Where do the poor, the prosperous and those in between live? Table 2 summarises. The richest 1 per cent are heavily concentrated in North America and Europe, where nine tenths of them live. This includes 15 per cent of the US population, 8 per cent of the UK population and 2 per cent of the entire European Union (EU) population. If we turn to the more inclusive global prosperous segment that is 3 per cent of the global population, then we find here 36 per cent of the US, 14 per cent of the UK and 8 per cent of the EU population —and 5 per cent of the population of Brazil. By contrast, among the global poor and insecure segments we find 90 per cent of the Chinese population, 60 per cent of Brazil and almost the entire populations of South Asia and sub-Saharan Africa, as opposed to just 12 per cent of the US and 13 per cent of the EU populations —and only 3 per cent 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 per cent of China's, but less than 1 per cent of India's population.

5.4 DIFFERENT PATTERNS OF GROWTH

These fractile charts can also help us identify different types of growth. For example, pro-poor growth, in which poor people benefit 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. Graphs for regions and countries are informative. We present a range of graphs in the Annex at the end of this paper and discuss below some of the more interesting insights they reveal. For example, the USA (Figure 39) is a classic example of consistently pro-rich growth across all consumption levels. In contrast, in the EU (Figure 40) growth has generally been relatively evenly spread across 90 per cent of the population, with the poorest 10 per cent, 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 41), 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 living on more than $10/day (Figure 42).

In contrast, in East Asia growth has been relatively distribution-neutral, with all segments benefiting more or less equally (see Figure 43 with China, and 44 excluding China).[30] China's growth, however, as is well documented, has been very strongly pro-rich (Figure 47). However, since even the lowest incomes have seen very large increases in per capita consumption, this has perhaps not yet been as destabilising socially as it might become. In a similar vein, India's growth is also pro-rich (Figure 48). The top 20 per cent have seen most of the benefits, and the top 5 per cent 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-elite but anti-poor and anti-middle too —meaning that the benefits of growth are actually negative at the lower end and concentrated at the top of the distribution (see Figure 45). In addition to the pro-elite growth pattern, four other 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 (Figure 46), South Africa, Malawi and Mali (see Annex). Second, there are examples of pro-middle growth, where those in the middle have benefited most —examples include Brazil (Figure 49). Third, there is anti-middle growth such as in Zambia (see Annex). Finally, there is surprisingly equitable growth across the distribution —this is the case in Vietnam (Figure 50), Nepal and the Philippines (see Annex).

FIGURE 39

USA

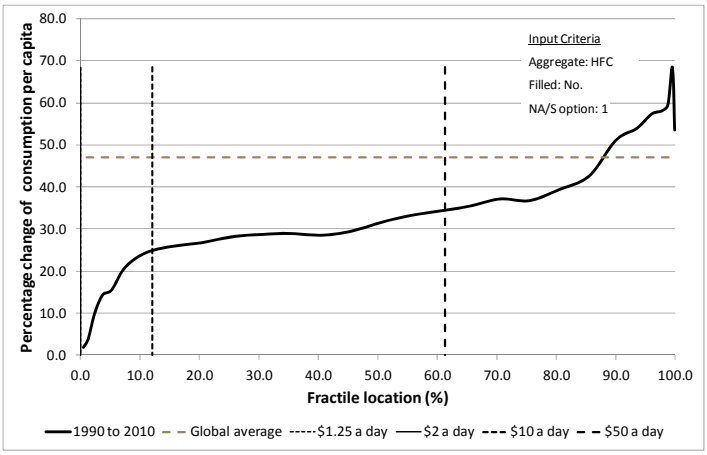

FIGURE 40

European Union

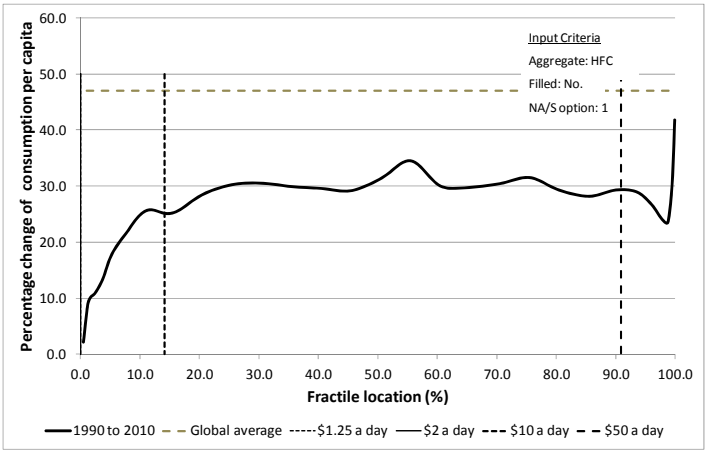

FIGURE 41

Latin America and Caribbean

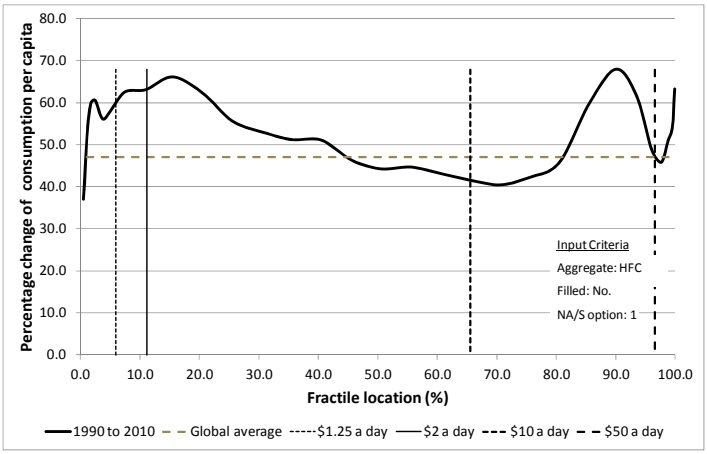

FIGURE 42

Latin America and Caribbean, Excluding Brazil

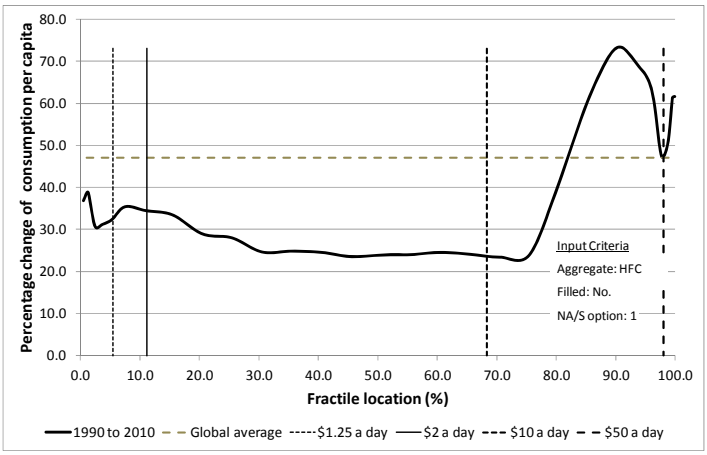

FIGURE 43

East Asia and Pacific

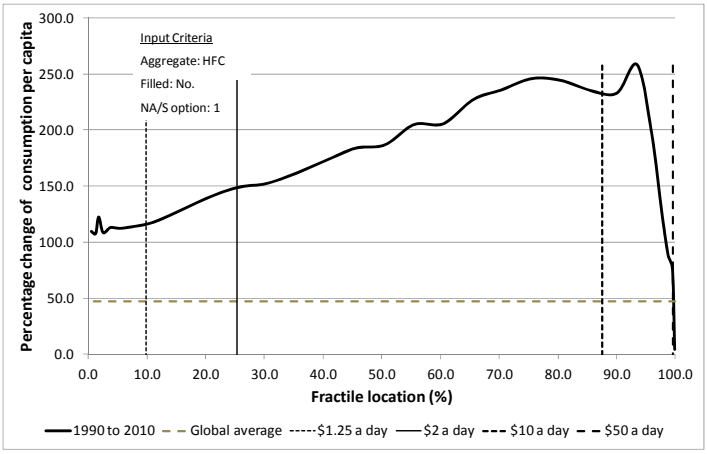

FIGURE 44

East Asia and Pacific, Excluding China

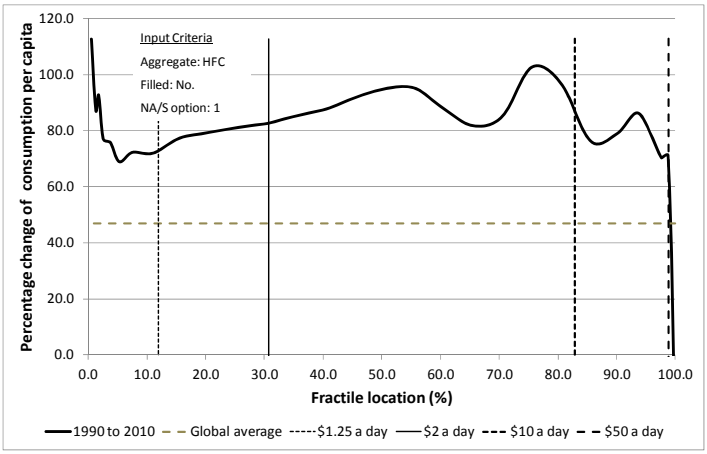

FIGURE 45

Nigeria – Anti-Poor and Middle and Pro-elite Growth

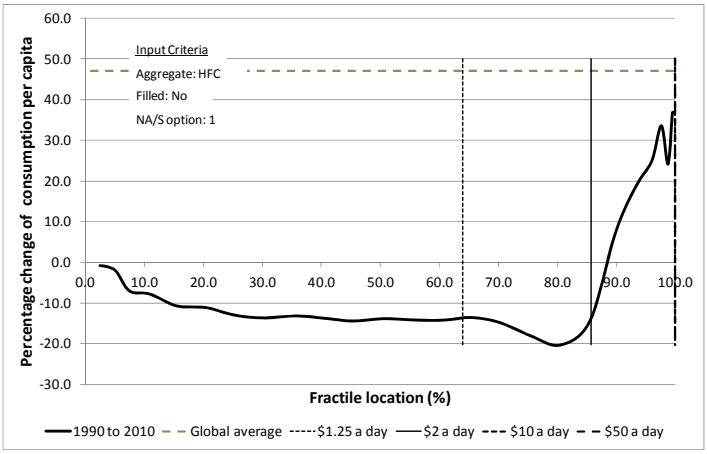

FIGURE 46

Ethiopia – Pro-poor Growth

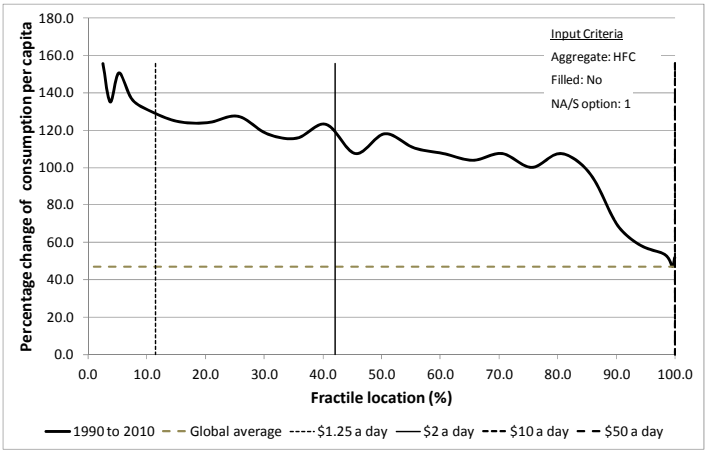

FIGURE 47

China

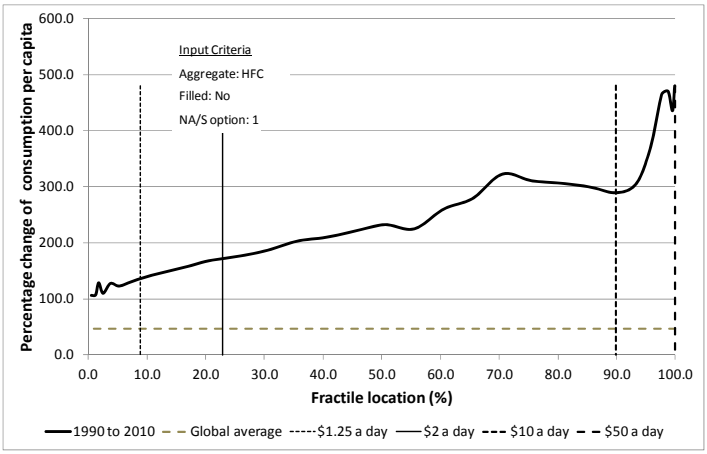

FIGURE 48

India

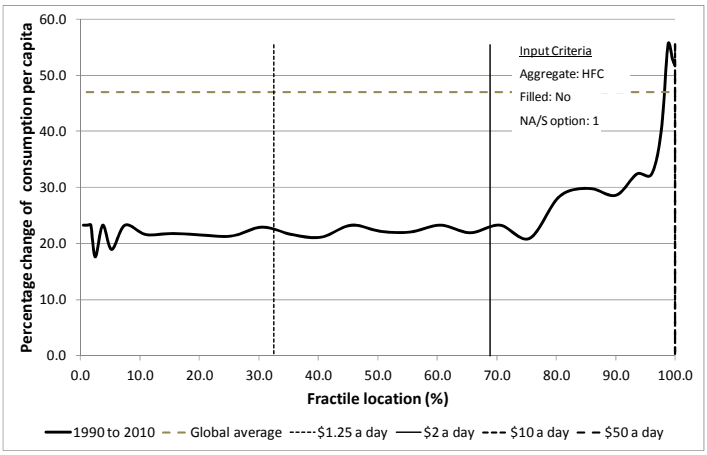

FIGURE 49

Brazil – Pro-middle Growth

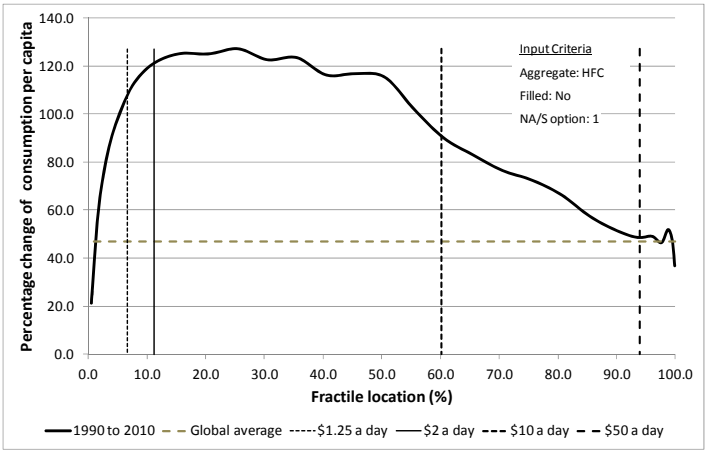

FIGURE 50

Vietnam – Equitable Growth

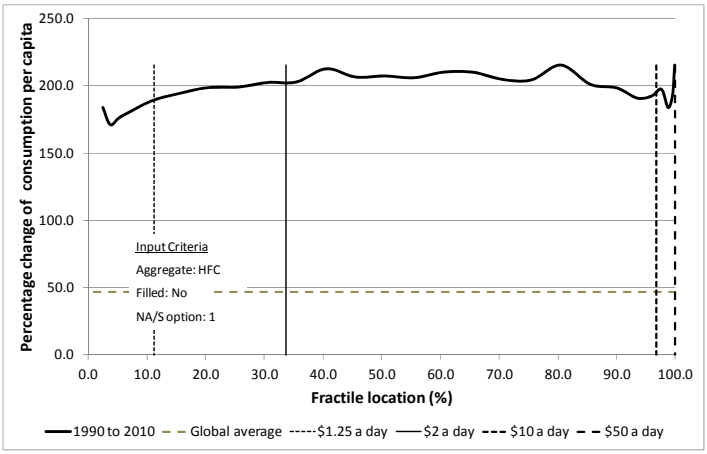

6 CONCLUSIONS

In this paper we have provided new estimates of the evolution of inequality between and within countries and focused on inequality in a multi-layered world in which we identify four 'layers': global absolute poor, global insecure, global secure and global prosperous. We have the following conclusions:

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, perhaps quite quickly in the last few years.

Not surprising, but little noted, is 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 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.

So 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/day); the 'global prosperous' (>$50/day) and those in between — specifically, the 'global insecure' ($2–10/day) and the 'global secure' ($10–50/day). One might also add a top layer of the 1 per cent of the population on over $75/day (equivalent to a family income of around $250,000 per annum 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 dumbbell) global consumption distribution. However, most of the 'structural' change in the distribution of global consumption is confined to UMICs. If we omit the UMICs, then the long-standing twin-peak dumbbell 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 15 per cent went to the richest 1 per cent of the global population. At the other end of the distribution, collectively, those living on under $2/day in 1990 (53 per cent of the population at that time) benefited from less than an eighth of the global growth from 1990 to 2010; and the 37 per cent of world population living on less than $1.25/day in 1990 collectively benefited by little more than a 20th of global growth.

And fourth, we identify five stylised growth patterns at the national level, as follows:

- pro-poor growth, by which those at the bottom end of the distribution have benefited most (e.g. Ethiopia);
- pro-middle growth, where those in the middle have benefited most (e.g. Brazil);
- anti-poor growth, meaning that the benefits of growth are negative at the lower end and concentrated at the top of the distribution (e.g. Nigeria);
- anti-middle growth (e.g. Zambia); and
- equitable growth (e.g. Vietnam), where the benefits are equal across the distribution.

This suggests that growth is not uniformly beneficial, and that is a point worth highlighting not least because we find that a modest amount of redistribution would have ended $2 poverty —if the share of global growth flowing to the 35 per cent of the global population who were living on $2/day in 2010 had increased from 5 per cent to 12 per cent (of that global growth over the 20 year period from 1990), this would have been sufficient to

end poverty at $2. In short, one could argue that the persistence of global poverty seems to have little to do with there being insufficient global growth and a lot more to do with a lack of collective will among the secure/prosperous cluster to forego a small share of their benefits from global growth in favour of a fairly modest amount of redistribution to the global absolute poor.

ANNEX

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							

TABLE A2

Estimates of each Segment of Global Population (millions) **by Region and Country Income Groups, 1990 and 2010: HFC, Survey Means** (Option 1)**, Filled, Coverage Compensated**

	Less than $2		$2–10		$10–50		$50+		Totals	
	1990	2010	1990	2010	1990	2010	1990	2010	1990	2010
Total	**2811**	**2407**	**1496**	**2910**	**905**	**1351**	**83**	**227**	**5296**	**6894**
East Asia (EAP)	1266	542	393	1267	148	370	14	23	1821	2202
Europe (ECA)	30	27	349	269	440	542	23	54	843	891
Latin America Caribbean	94	70	264	317	83	181	2	20	442	589
Middle East (MNA)	56	48	163	262	33	68	1	5	253	383
North America (NAM)	0	0	42	43	192	175	43	125	277	344
South Asia (SAR)	955	1092	191	537	0	4	0	0	1147	1633
Sub-Saharan Africa (SSA)	410	627	94	214	9	12	0	0	513	854
LICs	352	543	57	177	4	1	0	0	413	722
LMICs	1373	1459	443	1096	16	70	0	0	1832	2625
LMICs less India	653	616	289	716	16	68	0	0	958	1400
UMICs	1081	403	779	1488	201	486	2	27	2063	2405
UMICs less China	157	98	569	594	201	348	2	27	928	1068
HICs	4	1	217	148	684	793	82	199	987	1142
India	721	843	153	380	0	2	0	0	874	1225
China	925	305	211	894	0	138	0	0	1135	1338

TABLE A3

Estimates of Total Consumption of each Segment of Global Population ($billions per annum) **by Region and Country Income Groups, 1990 and 2010: HFC, Survey Means** (Option 1)**, Filled, Coverage Compensated**

	Less than $2		$2–10		$10–50		$50+		Totals	
	1990	2010	1990	2010	1990	2010	1990	2010	1990	2010
Total	**1063**	**1065**	**2509**	**4911**	**7338**	**10895**	**2341**	**6813**	**13251**	**23684**
East Asia (EAP)	467	273	540	2147	1076	2669	329	538	2411	5627
Europe (ECA)	17	13	746	613	3297	4532	506	1295	4566	6453
Latin America	39	30	511	643	648	1308	35	479	1233	2461
Middle East (MNA)	30	27	259	431	221	461	39	137	549	1055
North America (NAM)	0	0	98	98	2038	1802	1432	4360	3568	6260
South Asia (SAR)	394	521	230	680	0	19	0	0	624	1220
Sub-Saharan Africa	116	201	125	298	58	104	0	4	299	607
LICs	99	193	75	228	20	5	0	0	195	426
LMICs	552	669	580	1544	91	384	0	1	1223	2599
LMICs less India	241	269	392	1059	91	378	0	1	723	1708
UMICs	410	201	1325	2765	1379	3228	34	626	3147	6820
UMICs less China	71	46	1106	1214	1379	2436	34	626	2590	4321
HICs	2	1	530	374	5848	7277	2307	6186	8687	13839
India	311	400	188	485	0	6	0	0	499	891
China	338	155	219	1551	0	792	0	0	557	2498

SELECTED FRACTILE GROWTH CURVES (SURVEY MEANS, NOT FILLED – OPTION 1)

FIGURE A1

All HICs

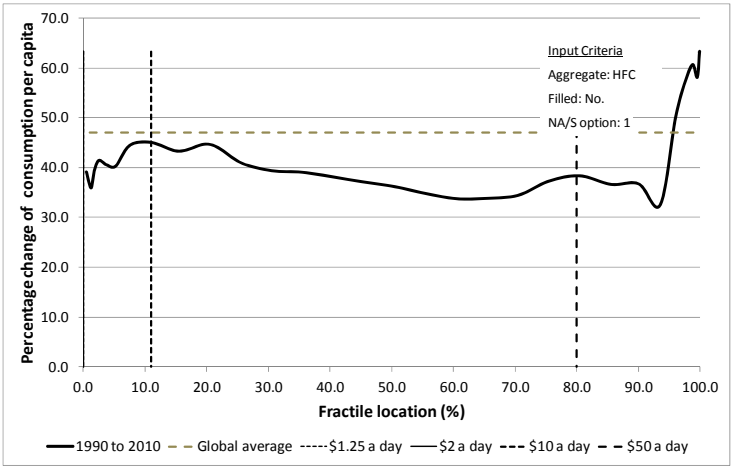

FIGURE A2

All UMICs

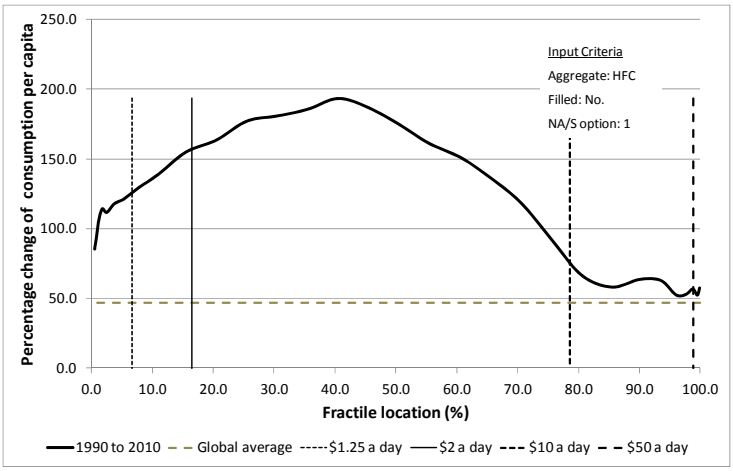

FIGURE A3

All UMICs, Excluding China

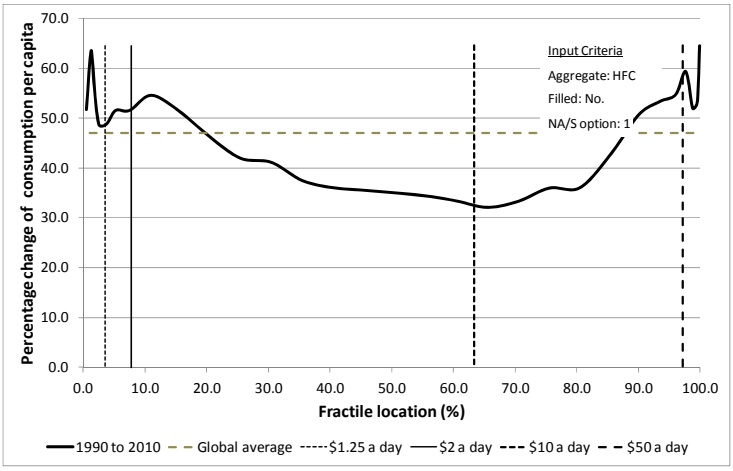

FIGURE A4

All LMICs

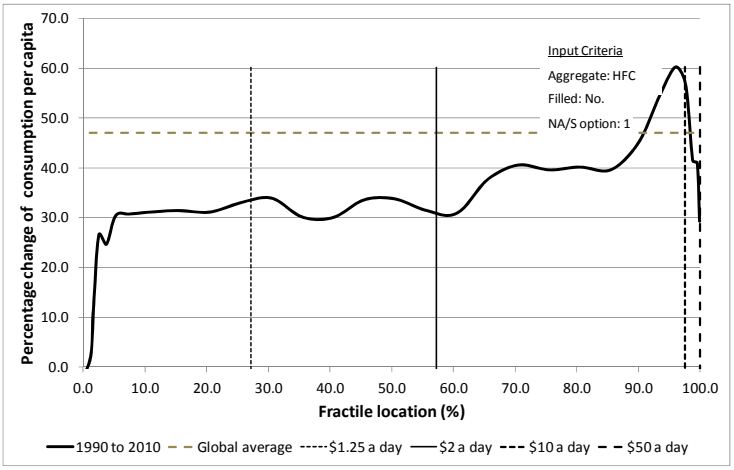

FIGURE A5

All LMICs, Excluding India

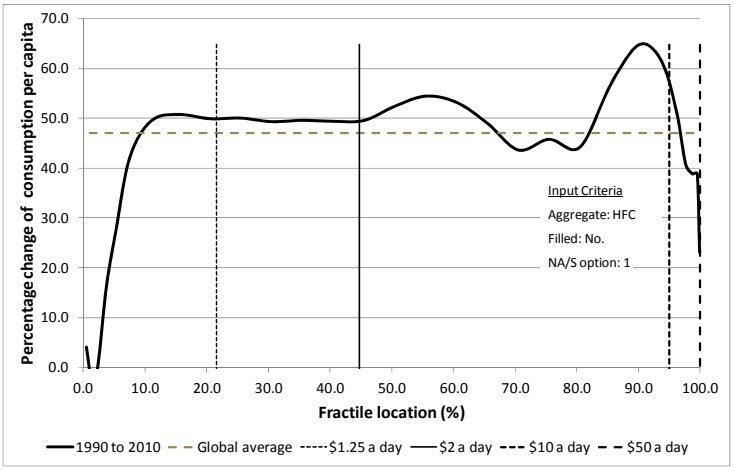

FIGURE A6

All LICs

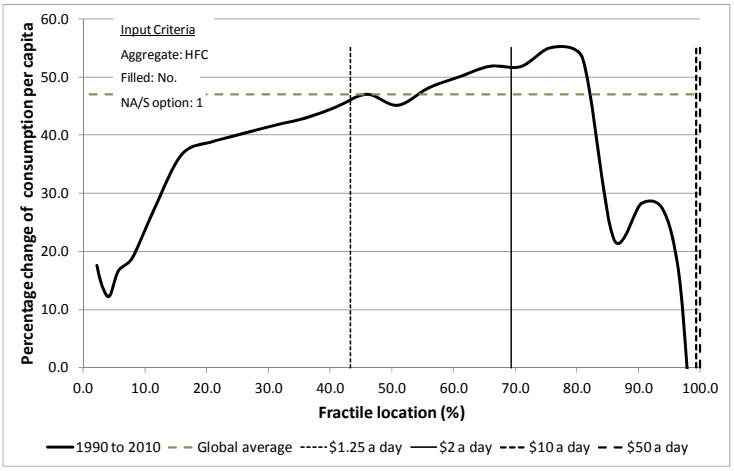

FIGURE A7

South Asia Region

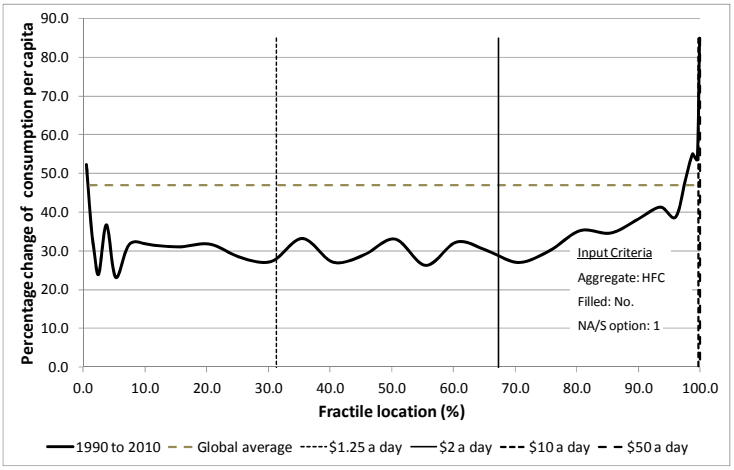

FIGURE A8

South Asia Region, Excluding India

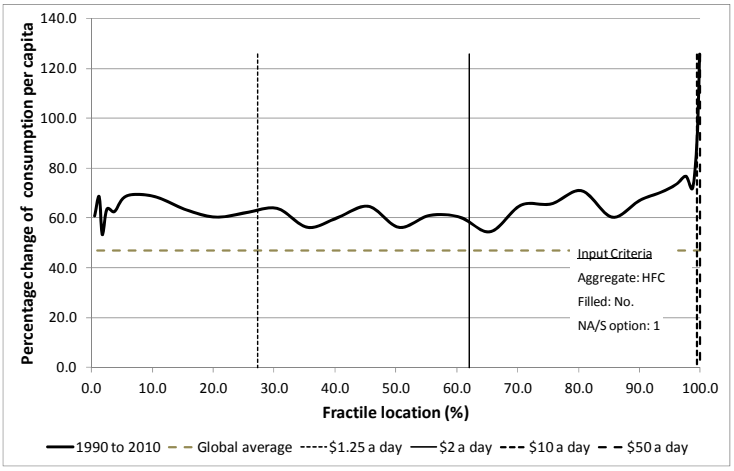

FIGURE A9

Sub-Saharan Africa

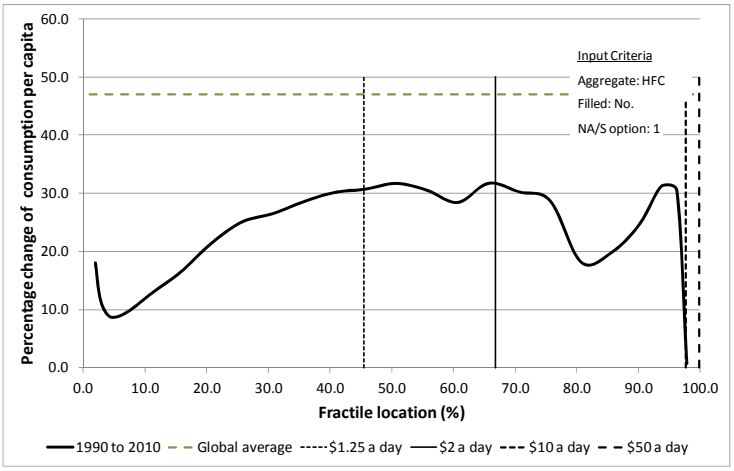

FIGURE A10

Sub-Saharan Africa, Excluding South Africa and Nigeria

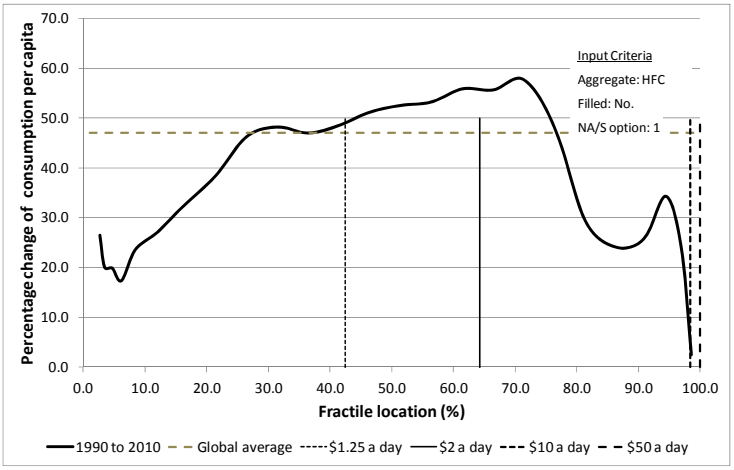

FIGURE A11

South Africa

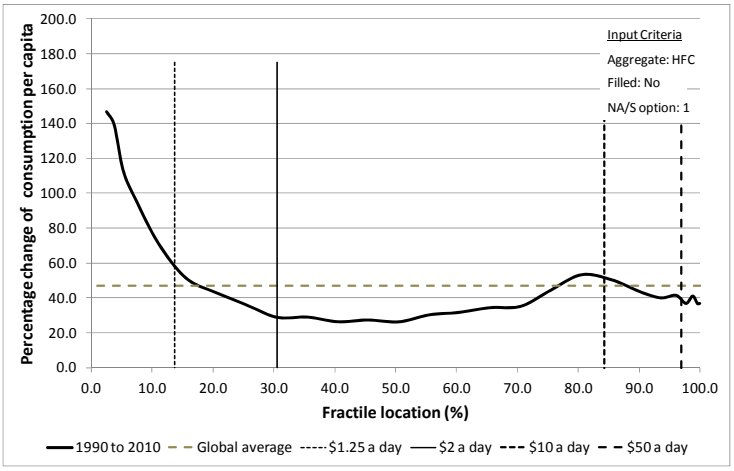

FIGURE A12

Bangladesh

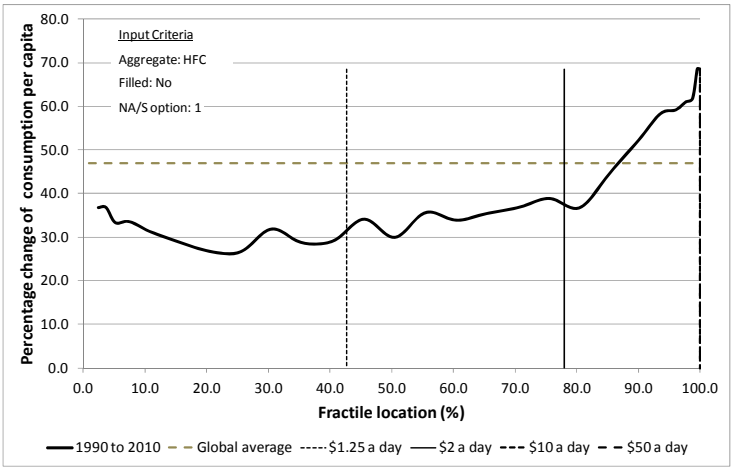

FIGURE A13

Indonesia

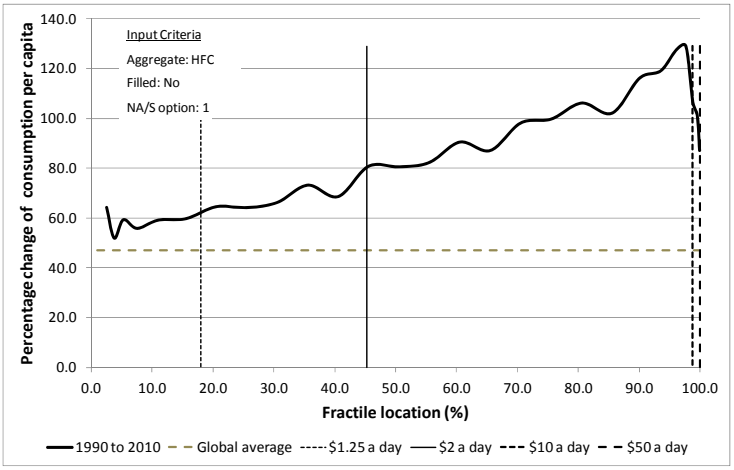

FIGURE A14

Malawi

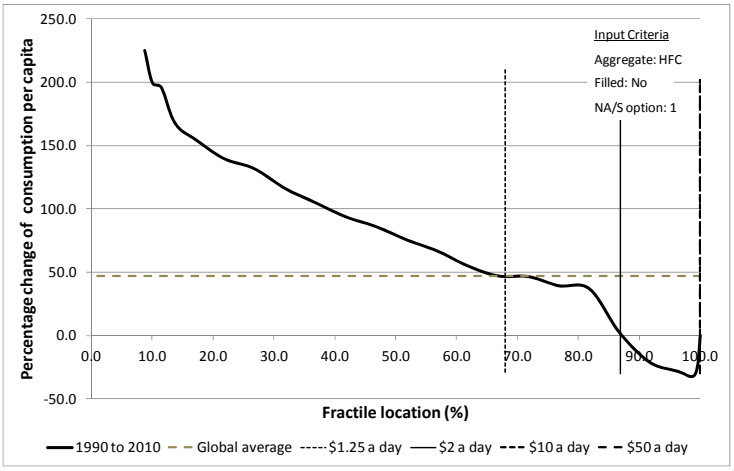

FIGURE A15

Mali

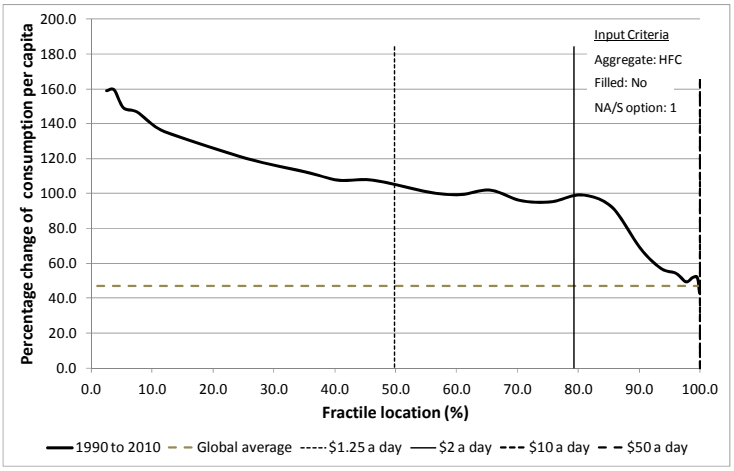

FIGURE A16

Mozambique

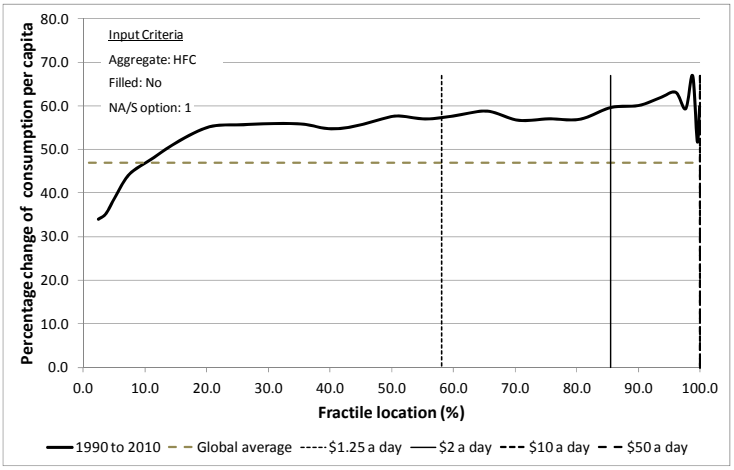

FIGURE A17

Pakistan

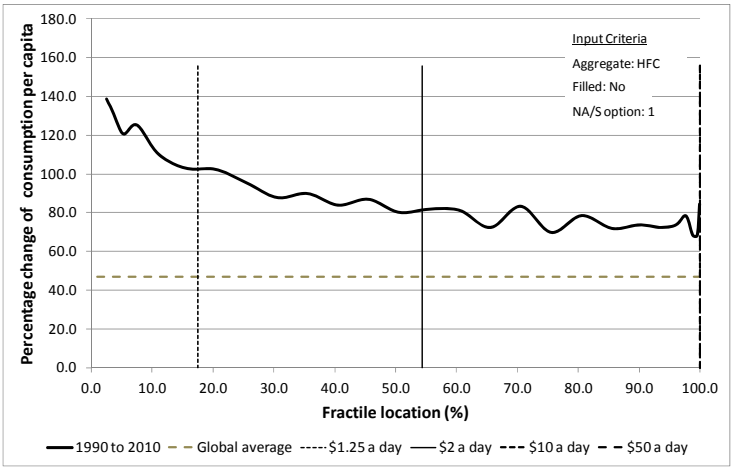

FIGURE A18

Zambia

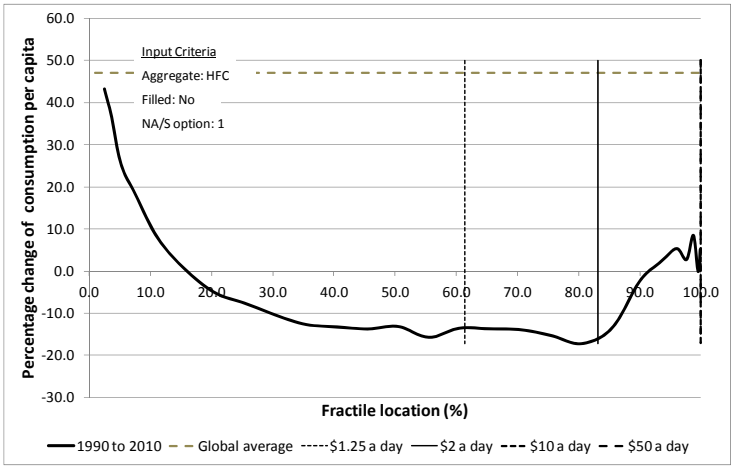

REFERENCES

Banerjee, E., and A.V. Duflo (2008). *What is Middle Class about the Middle Classes around the World?*. Boston, MA, Massachusetts Institute of Technology.

Bhalla, S. (2002). *Imagine there's no country: Poverty, inequality and growth in the era of globalization*. Washington, DC, Institute for International Economics.

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

Birdsall, N. (2010). 'The (Indispensable) Middle Class in Developing Countries; or the Rich and the Rest, Not the Poor and the Rest', *CGD Working Paper*, No. 207. Washington, DC, Centre for Global Development.

Bussolo, M., R. De Hoyos and D. Medvedev (2008). 'Economic Growth and Income Distribution: Linking Macroeconomic Models with Household Survey Data at the Global Level', *International Journal of Microsimulation*, 3(1): 92–103.

Carroll, W. K. (2012). 'Global, Transnational, Regional, National: The Need for Nuance in Theorizing Global Capitalism', *Critical Sociology*, 38(3): 365–371.

Chandy, L. and G. Gertz (2011). *Poverty in Numbers: The Changing State of Global Poverty from 2005 to 2015*. Washington, DC, Brookings Institution.

Chen, S. and M. Ravallion (2008). *The developing world is poorer than we thought, but no less successful in the fight against poverty*. Washington, DC, World Bank.

Chen, S. and M. Ravallion (2010). 'The Developing World is Poorer than we Thought, but no less Successful in the Fight Against Poverty', *Quarterly Journal of Economics*, 125(4): 1577–1625.

Chen, S. and M. Ravallion (2012). *An Update to the World Bank's Estimates of Consumption Poverty in the Developing World*. Washington, DC,World Bank.

Chotikapanich, D., D.S.P. Rao, . Griffiths and V. Valencia (2007). 'Global Inequality: Recent Evidence and Trends', *UNU-WIDER Research Paper*, Vol. 1: 1–32. Helsinki, United Nations University-World Institute for Development Economics Research.

Cobham, A. and A. Sumner (2013). *Putting The Gini Back In The Bottle? 'The Palma' As A Policy-Relevant Measure Of Inequality*. Mimeo.

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

Deaton, A. (2010). 'Price Indexes, Inequality, and the Measurement of World Poverty', *American Economic Review*, 100(1): 5–34.

Deaton, A. (2011). 'Measuring Development: Different Data, Different Conclusions', Paper presented at the 8th AFD-EUDN Conference. Paris, European Development Research Network.

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

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

Edward, P. (2006). 'Examining Inequality: Who really benefits from global growth?', *World Development*, 34(10): 1667–1695.

Edward, P. and A. Sumner (2013). 'The Future of Global Poverty in a Multi-Speed World', *CGD Working Paper*, No. 327. Washington, DC, Centre for Global Development.

Erikson, R., and J.H. Goldthorpe (1992). *The Constant Flux: A Study of Class Mobility in Industrial Societies*. Oxford, Clarendon Press.

Higgins, M. and J.G. Williamson (2002). 'Explaining inequality in the world round: Kuznets curves, cohort size, and openness', *Southeast Asian Studies*, 40(3): 269–288.

Hillebrand, E. (2008). 'The Global distribution of Income in 2050', *World Development*, 36(5): 727–740.

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

Kharas, H. and G. Gertz (2010). *New Global Middle Class: A Cross-Over from West to East*. Washington, DC, Brookings Institution Press.

Kharas, H. and A. Rogerson (2012). *Horizon 2025: Creative Destruction in the Aid Industry*. London, Overseas Development Institute.

Klasen, S. (2010). 'Levels and Trends in Absolute Poverty in the World: What We Know and What We Don't', Paper prepared for the International Association for Research in Income and Wealth, St. Gallen, Switzerland (22–28 August).

Lenagala, C. and R. Ram (2010). 'Growth elasticity of poverty: Estimates from new data', *International Journal of Social Economics*, 37(12): 923–932.

Linder, S.B. (1961). *An Essay on Trade and Transformation*. New York, NY, Wiley.

López-Calva, L. F. and E. Ortiz-Juarez (2011). 'A Vulnerability Approach to the Definition of the Middle Class', *Policy Research Working Paper*, No. 5902. Washington, DC, World Bank.

Lustig, N., Lopez-Calva, L., and 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, NJ, Princeton University Press.

Milanovic, B. (2011). 'Global inequality: from Class to Location, from Proletarians to Migrants', *World Bank Working Paper*, No. 5820. Washington, DC, World Bank.

Milanovic, B. (2012a). 'Global Income Inequality By the Numbers: In History and Now – an Overview', *World Bank Working Paper*, No. 6259. Washington, DC, World Bank.

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. and S. Yitzhaki (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*. Basingstoke, Palgrave MacMillan.

Palma, J.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', *Cambridge Working Papers in Economics*, No. 1111. Cambridge, University of Cambridge Department of Economics (later published in *Development and Change*, 42(1): 87–153).

Pinkovskiy, M. and X. Sala-i-Martin (2009). 'Parametric Estimations of the World Distribution of Income', *NBER Working Paper Series*, No. 15433. Cambridge, MA, National Bureau of Economic Research.

Ravallion, M. (2003). 'Measuring aggregate welfare in developing countries: How well do national accounts and surveys agree?', *The Review of Economics and Statistics*, 85(3): 645–652.

Ravallion, M. (2009). 'Do Poorer Countries Have Less Capacity For Redistribution?', *Policy Research Working Paper*, No. 5046. Washington, DC, World Bank.

Ravallion, M. (2010). 'The Developing World's Bulging (but Vulnerable) "Middle Class"', *Policy Research Working Paper*, No. 4816. Washington, DC, World Bank.

Ravallion, M. (2012). 'Benchmarking Global Poverty Reduction', *Policy Research Working Paper*, No. 6205. Washington, DC, World Bank.

Ravallion, M. and S. Chen (2012). *Monitoring Inequality*. Washington, DC, World Bank.

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*, No. w8933. Cambridge, MA, National Bureau of Economic Research.

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

UN Department of Economic and Social Affairs, Population Division (2011). *World Population Prospectus: The 2010 Revision*. New York, NY, UNDESA.

UNU-WIDER (2008). 'World Income Inequality Database'. Helsinki, United Nations University-World Institute for Development Economics Research.

World Bank (2007). *Global Economic Prospects: Managing the next wave of Globalization. Washington, DC, World Bank.*

World Bank (2012a). *'PovcalNet'. Washington, DC, World Bank.*

World Bank (2012b). World Development Indicators 2012. Washington, DC, World Bank.

NOTES

1. Elsewhere one could note the innovations in inequality research related to 'The World Top Incomes' database. See: <http://topincomes.g-mond.parisschoolofeconomics.eu.>.

2. The criteria for selection of articles here is (i) empirical in nature; (ii) long-term 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. In the text we only review studies since the 2008 update of PPPs (for 2005) because the PPP revision was substantial.

3. The criteria were the same as described in the above endnote.

4. See discussion in Cobham and Sumner (2013). Critiques of the Gini coefficient are not new, of course the Gini over emphasises how income is distributed and changes around the middle of the distribution.

5. Studies on Latin America argue that the changes in inequality are partly policy-induced (e.g. 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., 2013).

6. Nel (2012), for one, has noted that in much economic literature there is a tendency to disconnect studies of changes in within-country inequality from the broader economic integration of countries into contemporary globalisation and global production patterns. Many scholars have considered the role of globalisation in contributing to within-country inequality. Not least are those who focus on the functional distribution of income.

7. Ravallion and Chen (2012) have noted somewhat similar findings that "inequality within growing developing countries falls about as often as it rises ... The evidence leads one to doubt that higher inequality is simply the 'price' for higher growth and lower absolute poverty" (2012b: 5). The issue with this and other studies is that it is not really a question of whether 'growth' improves or worsens inequality. It is what specific macroeconomic policies have what effect. For example, some types of macro policy are equity-enhancing, and some are not, and there are examples of both profit-led and wage-led growth.

8. Cobham and Sumner (2013) corroborate and explore Palma's (2011) findings.

9. 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 per cent (in percentage points) minus 15, which implies (if one assumes that the share of the lowest 40 per cent 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 per cent (2011: 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). Palma also explores whether there are any statistical relationships between inequality and mean national income (GDP) per capita, finding that most low- and middle-income countries/regions have Ginis of around 40, whereas it is mainly among the rich countries that the greatest distributional diversity is found – ranging from a Gini of about 46 in the USA down to close to 25 in Japan and the Nordic countries (data from WDI). 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.

10. The dataset is publically available at <http://econ.worldbank.org/projects/inequality>.

11. Inequality 1 is "focused on inequality between nations of the world. It is an inequality statistic calculated across GDPs or mean incomes obtained from household surveys of all countries in the world, without population-weighting" (Milanovic, 2012: 3). In short, 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. Milanovic argues that the "mother of all inequality disputes" is due to Inequality 1 and 2 moving in different directions between 1950 and 2010, with Inequality 1 (not population-weighted) rising, and Inequality 2 (population-weighted) falling (2012: 6). 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).

12. From Milanovic's Figure 4 (2012: 13), this 'class' appears to be the global population between the 25th and 60th percentiles.

13. We note that one issue with 'Inequality 3' models (and this includes GrIP) is that it is often national-level policies that matter in alleviating inequality, so that identifying those countries where inequality is rising or falling may be as relevant or more so than attention to global inequality trends across individuals. We address differences in national inequality trends at the end of this paper.

14. While Milanovic does not define explicitly the 'global middle class', Ravallion (2010) does do so by per capita expenditures (as does Kharas, 2010a; 2010b). 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: 446).

15. Ravallion and Chen (2012: 4) also note significant variations between regions. For example, the region with the highest average inequality within countries is Latin America and the Caribbean (LAC), in which 90 per cent of inequality is within countries, although this has been falling since 2000. The second-highest average within-country inequality is in sub-Saharan Africa (SSA). And South Asia and East Asia both have generally had low within-country inequality, but this has been rising since the 1990s.

16. See: <www.kcl.ac.uk/aboutkings/worldwide/initiatives/global/intdev/people/Sumner/Edward-Sumner-Version04March2013.pdf>.

17. In this paper we use the Povcal version of February 2012.

18. 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>.

19. 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 is difficult to adjust for in subsequent analysis. As is the case for most other analysts 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.

20. In the GrIP model between-country inequality is calculated by assuming there is no within-country inequality (i.e. all members of a national population are deemed to have the same consumption per capita). Within-country inequality is derived by assuming that all countries have the same average income (so that the only differences in consumption arise in the model from intra-national inequality).

21. Milanovic estimates that in 1870 the global Theil index was about 65, with two thirds of global inequality being due to within-country inequality – so the world still has a long way to go before we get back to that situation.

22. EAP = East Asia and Pacific; LAC = Latin America and Caribbean; SAR = South Asia Region; SSA = sub-Saharan Africa. High Inc = All high-income countries in the world.

23. As alluded to in the earlier discussion, a body of empirical studies related to developing countries has emerged in response to the growing data on the 'in-between' groups, often referred to as the non-poor/non-rich or the 'non-polar' groups or classes. Typically referring to these 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 the review by Sumner, 2012). Some have taken a relative approach by defining either the literal 'middle' of the income/expenditure distribution in terms of the middle three expenditure quintiles, or the non-literal middle as those between the 'poor' (taken as the bottom 40 per cent) and the 'rich' (taken as the top 10 per cent) (e.g. Palma, 2011).

24. 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'. It would, we fear, be too readily misunderstood and misquoted, so we prefer to present these curves in the dimensionless form used. That approach also allows us to present the population and consumption curves in one graph on the same scale.

25. The old twin-peaks world was identified by Quah (1997). The likelihood of the trend away from a two-peak world was previously detected in Edward (2006: 1677). As noted earlier, the recent 'consumption' density curves also show the emergence of a new peak at the highest consumption levels, i.e. approaching $100,000 PPP per capita per annum. This was present back in 1985 but seems to be becoming more distinct recently – perhaps indicative of the emergence of a 'super-rich' segment pulling away from the global prosperous. However, this phenomenon needs to be treated circumspectly for reasons discussed in the text.

26. We did review the equivalent curves using filled data to detect whether there were any significant differences, and in most cases the filled analysis supported the general trends and findings from the unfilled graphs, although usually with slight differences in absolute values.

27. Milanovic (2012: 13) has recently presented a similar figure for the period from 1988 to 2008, while an alternative version of the same sort of chart was previously published by Edward (2006: 1681). Milanovic estimated that the top 1 per cent globally saw its real income rise by 60 per cent in the two decades from 1988. Those between the 75th and 90th percentiles saw a rise of just 10 per cent 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 per cent and 80 per cent. 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 period from 1988 to 2008 that Milanovic used, and our findings there were similar to those we describe here for 1990 to 2010. The main differences from Milanovic's results are that in our results

the main winners (other than the richest few per cent) appear to be between the 45th and 75th percentiles, whose incomes rose by between 60 per cent and (almost) 90 per cent – a smaller group than Milanovic found. And although 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.

28. It is worth reiterating at this point both that this analysis mostly excludes the global super-rich, because their income is rarely fully captured in national surveys, and that the method of estimation used in GrIP tends to be conservative about the incomes of the very richest 1 to 2 per cent in each country. The $250,000 income figure is derived as follows. In GrIP, the top 1 per cent are those with individual consumption levels at or above $75/day (in 2005 $PPP) in 2010. If we envisage a typical family of four, then that would imply a household consumption of $300/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 per cent 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/day equivalent to a household income of about $250,000 a year – a far from insignificant income level even in the developed world today.

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

30. 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 per cent of the distribution) of the graphs – as occurs, for example, in the curve for EAP without China above the $50/day level. At these ends the analysis can be particularly sensitive to slight changes in the data, so these large changes may not be reliable representations of actuality.

International Policy Centre for Inclusive Growth (IPC - IG)
United Nations Development Programme

SBS, Quadra 1, Bloco J, Ed. BNDES, 13º andar
70076-900 Brasilia, DF - Brazil
Telephone: +55 61 2105 5000

E-mail: ipc@ipc-undp.org ▪ URL: www.ipc-undp.org