

Measuring Global Poverty: Why PPP Methods Matter

Robert Ackland* Steve Dowrick† Benoit Freyens‡

September 17, 2008

Abstract

We present theory and evidence to suggest that, in the context of analysing global poverty, the EKS approach to estimating purchasing power parities yields more appropriate international comparison of real incomes than the Geary-Khamis approach. Our analysis of the 1996 International Comparison Project data confirms that the Geary-Khamis approach leads to substantial overstatement of the relative incomes of the world's poorest nations and to misleading comparisons of poverty rates across regions. Similar bias is found in the Penn World Table which uses a modified version of the Geary-Khamis approach. Estimation of both the level of global poverty and its location is very sensitive to the choice of index. The EKS index of real income is much closer to being a true index of economic welfare and is therefore to be preferred for assessment of global poverty.

*Australian Demographic and Social Research Institute, The Australian National University, robert.ackland@anu.edu.au

†Corresponding Author. School of Economics, The Australian National University. Postal address: Arndt Building (no. 25A), Kingsley Street, The Australian National University, ACT 0200, AUSTRALIA. TEL: +61 02 6125 4606 ; FAX: +61 02 6125 5124, steve.dowrick@anu.edu.au

‡School of Economics and Research School of Social Sciences, The Australian National University, ben.freyens@anu.edu.au

1 Introduction

A halving of extreme poverty by 2015 is the first of the United Nation's Millennium Development Goals and approaches for estimating world poverty have been widely discussed in recent years by authors such as Atkinson and Brandolini (2001), Bhalla (2004), Chen and Ravallion (2001), Deaton (2001, 2005), Ravallion (2003a,b), Reddy and Pogge (2003) and Sala-i-Martin (2006). This paper draws on insights from the economic approach to index numbers to highlight the fact, little appreciated in most of the empirical work on global poverty, that global estimates of the number of people living in poverty and their distribution across countries and regions depend crucially on the purchasing power parities that are used to translate a common poverty line into local currencies.

In particular, we draw attention to the significant differences that arise from using the EKS index method rather than the Geary-Khamis¹ index on which both the Penn World Table (Heston, Summers, and Aten, 2002) and Angus Maddison (Maddison, 1995, 2003) base their measures of real GDP. Since 2000, the World Bank's purchasing power parity estimates used for measuring global poverty and inequality have no longer been based on the Geary-Khamis method, but instead constructed with the EKS method (Chen and Ravallion, 2007). Bhalla (2002, 2004) and Sala-i-Martin (2002, 2006) discuss some of the problems that arise in the measurement of global poverty, and provide significantly lower estimates than those presented in World Bank (2000, 2002), but fail to question the use of the Geary-Khamis method in the Penn World Table data that they rely on. Nor do Berry, Bourguignon, and Morrisson (1983) and Bourguignon and Morrisson (2002) question the Geary-Khamis method underlying the income data that they take from Maddison (1995) in their analysis of the shifting composition of the world's poor.

In this paper we demonstrate and explain the importance of the choice between alternative methods of calculating purchasing power parities. Our analysis suggests that the EKS method produces income measures which are closer to economic definitions of true welfare. We find that the Geary-Khamis method understates the number of the world's poor by 30 percent relative to the EKS method. Moreover, the Geary-Khamis method understates

substantially the proportion of the world's extremely poor population who live in Asia.

It is important to note that we do not attempt to provide a definitive count of the world's poor, nor to offer judgements on disputes between the World Bank and its critics on issues such as the choice between national income or survey measures of average income. Our analysis is intended to quantify the orders of magnitude involved in getting the purchasing power parity concept and method right.

In Section 2 we explain the difference between the Geary-Khamis and the EKS methods of calculating purchasing power parities. We explain why the Geary-Khamis method may exaggerate the purchasing power of poor country currencies and we discuss the properties of the EKS index.

In our third section we apply the two purchasing power parity methodologies to the latest International Comparison Project data set covering 115 countries. We also introduce the notion of "true income" comparisons based on the economic theory of index numbers that has been developed by Afriat (1967, 1984), Varian (1983) and Dowrick and Quiggin (1997), enabling us to quantify the direction and magnitude of bias imparted by the two methods. Our analysis confirms the findings of Hill (2000), who reports that the Geary-Khamis method imparts biases up to a factor of 2 in bilateral international income comparisons relative to the Fisher index, and the findings of Dowrick and Akmal (2005) who report that the Geary-Khamis method understates true global income inequality relative to the multilateral Afriat index.

This paper goes well beyond both of these papers in that, having established the direction and magnitude of bias in PPP methods, we then demonstrate the implications of this bias for global poverty estimates. In Section 4 we discuss alternative approaches to measuring world poverty and explain our choice, which involves estimating the cumulative income distribution function for 95 of the International Comparison Project countries using published data on quintile shares. We are then able to demonstrate, in Section 5, the impact of changing from Geary-Khamis to EKS measures of purchasing power parity in terms of estimating both the total number of people in poverty and the distribution of the

poor population across and within regions.

2 An explanation of the different purchasing power parity methods

It is well known that international currency markets tend to undervalue the domestic purchasing power of currencies of low productivity / low income countries – the Balassa-Samuelson effect (Balassa, 1964; Samuelson, 1964). Real wages are low in countries with low labour productivity, so non-traded labour-intensive services are cheap relative to traded goods in poorer countries. Market exchange rates tend to equate the prices of tradables across countries. This causes the foreign exchange market to undervalue the domestic purchasing power of the currencies of poor countries, hence to understate their real incomes relative to income levels in rich countries.

Alternatives to exchange-rate comparisons of income rely on the estimation of purchasing power parities (PPP). A massive research effort, based on detailed price surveys in many countries under the auspices of the International Comparison Program (ICP), has resulted in the publication of the Penn World Table (PWT) ² which provides measures of *real* GDP per capita at constant international prices for over one hundred countries since 1950. These data are commonly referred to as PPP measures of real income.

Many users of the PWT data are unaware, however, that attempts to measure purchasing power are problematic. The PWT estimates of average real income (GDP *per capita*) are based on the Geary-Khamis (GK) method of constructing “average international prices”. The GDP of each country is evaluated at these fixed prices. This method does succeed in removing the traded-sector bias in exchange rate valuations, yielding real income differentials which are substantially smaller than the income differentials of exchange rate comparisons. The use of fixed price valuations, however, introduces systematic bias by ignoring substitution towards higher consumption of goods and services that are locally

cheap. Evaluated at non-local prices, this higher consumption is misinterpreted as higher income. The use of a rich country's prices to conduct an income comparison will tend to overstate a poor country's relative income, and vice versa.

Formally, GK measures of real GDP and the corresponding purchasing power parities are derived as the solution to the following set of simultaneous equations:

$$\begin{aligned}\pi_k &= \sum_i \frac{p_k^i}{e^i} \frac{q_k^i}{\sum_i q_k^i} \quad k = 1 \dots K \text{ items of expenditure} \\ e^i &= \frac{\mathbf{p}^i \cdot \mathbf{q}^i}{\boldsymbol{\pi} \cdot \mathbf{q}^i} \quad i = 1 \dots N \text{ countries}\end{aligned}\tag{1}$$

The observed national price and quantity vectors for country i are represented by \mathbf{p}^i and \mathbf{q}^i respectively. $\boldsymbol{\pi}$ is the imputed international price vector, comprised of the prices for individual commodities, π_k . The imputed purchasing power parity, e^i , is defined so that each country's GDP bundle valued at international prices, $\boldsymbol{\pi} \cdot \mathbf{q}^i$, is equal to its local currency GDP deflated by the purchasing power parity, $\mathbf{p}^i \cdot \mathbf{q}^i / e^i$.

The direction of bias imparted by this method is to underestimate the true income differentials between rich and poor countries. This is because the GK method calculates international prices as the weighted average of national price vectors, where the weights on each item's prices are each country's share of global consumption on that item. This approach ensures that the price relativities in the international price vector are closer to the relativities of the richest and most populous of the OECD countries than to those of the world's poorer economies, given that 45% of world GDP in 1996 was produced by just 7 OECD countries –USA, Japan, Germany, France, UK, Italy and Canada.

Substitution bias implies that using rich country prices to value GDP in every country tends to overstate the relative income levels of poor countries. Use of GK estimates of real national incomes avoids the bias which arises from the use of market rates of exchange, but it introduces an opposite bias which tends to understate the true extent of world poverty.³

An alternative way of comparing real incomes across countries is to use the EKS index,

based on the bilateral Fisher index, F^{ij} which is constructed as the geometric mean of the Paasche and Laspeyres indices:

$$P^{ij} = \frac{\mathbf{p}^i \cdot \mathbf{q}^i}{\mathbf{p}^i \cdot \mathbf{q}^j}; \quad L^{ij} = \frac{\mathbf{p}^j \cdot \mathbf{q}^i}{\mathbf{p}^j \cdot \mathbf{q}^j}; \quad F^{ij} = (P^{ij} L^{ij})^{0.5} \quad (2)$$

Substitution bias implies that the Paasche index is liable to understate the true income ratio of the richer country relative to the poorer country whilst the Laspeyres index is liable to overstate it. It is not necessarily the case that the two biases will exactly cancel out in the Fisher index, but nor is there any expectation of systematic bias. The expectation of no systematic bias carries over to the multilateral EKS index which is constructed as the geometric mean of Fisher indices:

$$EKS^{ij} = \prod_{n=1}^N \left(\frac{F^{in}}{F^{jn}} \right)^{1/N} \quad (3)$$

The OECD currently uses the EKS method for comparisons of incomes between their member countries. We need to construct the EKS indices ourselves from ICP data as any investigation into global poverty must extend its country coverage well beyond the OECD.

3 Comparing GK, EKS and True Income Levels

We use the most recent survey data from the ICP covering 115 national price surveys in 1996.⁴ The data consist of prices and *per capita* expenditures in local currency on 31 items which sum to GDP per capita, based on the familiar national accounting identity: $Y = C + I + G + NX$. From these data we calculate the real *per capita* quantities of each item in each of the 115 countries surveyed. Given prices and quantities, we can calculate both the GK and EKS measures of real GDP per capita. We are also able to calculate the true Afriat index for a sub-sample of countries, as explained later.

We construct the GK index by solving the system of equations (1), normalising the parities so that $e^{\text{USA}} = 1$, giving a vector of national purchasing power parities, \mathbf{E}^{GK} , and a vector

of real GDP *per capita*, y^{GK} , where the value of real GDP for the USA equals its local currency value.

The EKS index is constructed from (2) and (3) and we normalise this to the index y^{EKS} , where the value for the USA equals its local currency GDP. The 1996 ICP did not include either China or India, omitting a third of the world's total population and a substantial proportion of the world's poor. Although we are unable to directly calculate the GK and EKS income measures for these countries, we do have the PWT estimates of real GDP *per capita*, which have been derived by extrapolation from the ICP sample and from other surveys. We are able to estimate a within-sample relationship between PWT and EKS incomes which allows us to estimate the EKS income level for both China and India, giving us our 'extended ICP' sample of 117 countries. This extends our sample of countries for which we have EKS income estimates to include 91% of the world's population and 95% of the world's GDP.

Using this extended sample we find that the two measures of log GDP are highly correlated. The scatter plot is shown in Figure 1 with log GK income on the vertical axis and log EKS income on the horizontal axis. The 45 degree line and the following OLS regression line are also shown (note that the standard error of coefficient estimates is in brackets):

$$\ln y_i^{GK} = 0.679 + 0.940 \ln y_i^{EKS} \quad R^2 = 0.990, \quad s.e. = 0.103, \quad n = 115 \quad (4)$$

(0.077) (0.009)

The slope coefficient of 0.94 is significantly less than unity; this, combined with the low standard error of the regression (0.103) implies that, as expected, the GK measure tends to compress the distribution of income across countries relative to the EKS measure.⁵ For a poor country such as Yemen, with real income one fiftieth that of the USA, the regression coefficient implies that the GK method tends to overstate the real income ratio by more than one quarter relative to the EKS method, as illustrated by the vertical gap between the regression and 45 degree lines in Figure 1.⁶

3.1 Evaluating Bias in the GK and EKS Indexes

Evidence that the GK and EKS income ratios differ substantially suggests that the choice of method is likely to have a significant influence on estimates of the level of global poverty and its distribution across countries. This raises the question of which is the preferred method of comparing incomes at purchasing power parity: whilst we expect the GK income ratios to be biased upwards, how confident are we that the EKS measures are unbiased? There are reasons to suspect that the EKS index is relatively free from bias because the substitution bias inherent in the Laspeyres index is likely to be offset by the opposite bias in the Paasche index. It is not clear, however, whether the biases cancel out exactly.

To address these questions we make use of the theory of consumer behaviour which underlies the economic approach to welfare indexes. In the context of national income or GDP we suppose the existence of a representative national household with preferences over the inter-temporal consumption bundle, leading to rational choices over current consumption and investment. In order to make meaningful international comparisons we further suppose that representative households have common preferences. If these assumptions are not rejected by the data, we can use the money-metric Allen welfare index which compares the minimum expenditures required to achieve utilities U^2 and U^1 at some reference price vector, \mathbf{p}^r :

$$A_r^{2:1} = \frac{e[U(\mathbf{q}^2), \mathbf{p}^r]}{e[U(\mathbf{q}^1), \mathbf{p}^r]}$$

The Allen index is, in general, dependent on the choice of reference price vector, unless preferences are homothetic. If a set of price and quantity data can be shown to be consistent with common homothetic preferences, we can use theorems from Afriat (1967, 1984) to determine tight bounds on the Allen index, which we refer to as the True Afriat index. We can then define bias as the extent to which either the GK or EKS indexes violate the true bounds.

Typically, studies of purchasing power parity do not test the hypothesis of common preferences, nor do they know the form of the preference relationship. Here we report on

tests for common preferences and on tests for homothetic preferences, noting that linear quadratic functions, for which the EKS index is exact, comprise a subset of the more general homothetic set. We follow Dowrick and Quiggin (1994) in using revealed preference relationships to test for common preferences. The hypothesis of common preferences is rejected for countries A and B if we observe that the Laspeyres index is strictly less than one and the Paasche index is strictly greater than one, i.e. A could have afforded B's commodity bundle and B could have afforded A's bundle. Applying these tests to the 6,555 bilateral comparisons amongst the 115 ICP countries, we find only one pair where common preferences are rejected: Armenia and Uzbekistan.

The test for common homothetic preferences is much stricter. Afriat (1984) has shown that the test requires that there exist a true multilateral index, \mathbf{a} , such that the ratio a_i/a_j , lies between the Paasche and Laspeyres indexes for every pair of countries i and j . We use Varian's(1983) minimum path algorithm to identify a set of 80 countries that satisfy the test. The fact that nearly one third of the ICP countries do not satisfy the test is a major weakness in applying the Afriat approach to constructing a comprehensive multilateral index. Nevertheless, we can use the Afriat results to assess the degree of bias in the GK and EKS indexes within the set of countries for which aggregate homotheticity is not rejected.

Satisfying the Afriat test does not imply a unique Afriat index. Rather, there are well-defined bounds. Within these bounds there is an irreducible indeterminacy resulting from the fact that we do not observe utility. We make use of results from Dowrick and Quiggin (1997) to assess the extent to which the EKS and GK measures violate the bounds, $a_i^- \leq I_i/\bar{I} \leq a_i^+$, where \bar{I} is the geometric mean of the index I which is being tested. The bounds \mathbf{a}^- and \mathbf{a}^+ are constructed following the method described in Dowrick and Quiggin (1997). They represent the lower and upper bounds of true indices relative to the mean.

The results are illustrated in Figure 2 where the solid lines represent the upper and lower Afriat bounds, \mathbf{a}^- and \mathbf{a}^+ , and the GK and EKS indexes are displayed as open and closed boxes respectively. Both indexes are in logs and have been normalised to a zero mean. The countries are ordered from the poorest, Tanzania, to the richest, USA, along the horizontal

axis.

In Figure 3 we add the PWT6.1 measures of real GDP *per capita* to our comparison. To enable closer inspection we display each index in terms of its deviation from the mid-point of the Afriat bounds. We find that both the PWT and the GK indexes violate the bounds for around two-thirds of the observations (54 and 51 cases out of 80, respectively) and the mean absolute deviation outside the bounds is close to 11% for both measures. For both PWT and GK, there is a very obvious overestimation of real GDP for the poorest countries and underestimation for the richest.

On the other hand, when we examine the deviations of the EKS measures we find that 80% of the observations lie within the bounds. The 16 deviations outside the bounds are small, averaging only 5%. Furthermore, there is no tendency for the EKS to systematically under- or over-estimate the true GDP of the relatively rich or relatively poor.

Given the absence of systematic bias in the EKS index relative to the true Afriat bounds and the evidence of substantial and systematic bias in both the GK and PWT indexes for our sub-sample of 80 countries, we regard the EKS index as far more suitable for the measurement of world poverty. We note that the PWT income measures are likely to substantially overstate true income levels in the world's poorest countries relative to the USA; hence any poverty measurement based on the PWT is likely to substantially understate the extent of poverty when the international poverty line is defined in US dollars. Moreover, since the degree of bias in PWT income varies from country to country, poverty estimates based on the PWT are likely to misallocate the world's poor across countries.

We use the bounds of the Afriat index as a benchmark against which to measure bias in other indexes because it is, by construction, a utility-based measure which is free of substitution bias. It has an explicit money-metric interpretation in terms of the welfare of a representative consumer facing the average income and the prices of each country. The representative consumer is, of course, a convenient fiction; individual consumers may well have heterogeneous and non-homothetic preferences. We are, however, able to exploit the remarkable finding that aggregate (per capita) behaviour in a majority of countries does not

contradict the hypothesis that such a representative consumer does choose each country's GDP bundle when faced with that country's budget constraint.

4 Approaches for estimating world poverty

A first step in measuring global poverty is the setting of an International Poverty Line (IPL), defined in terms of an absolute level of consumption or income, on either an individual or household basis. The World Bank's '\$1/day/person' and '\$2/day/per person' IPLs (see, for example, World Bank, 2000) provide the yardstick by which most global poverty comparisons are produced today.

There are two main approaches that have been used for estimating the total number of people living below the IPL, or the "absolute poor". World Bank estimates of global poverty (see, for example, Chen and Ravallion, 2001) have involved first converting the IPL into local currencies using PPP exchange rates and then using nationally-representative surveys of household consumption to estimate the number of people living below these "national poverty lines". An alternative approach, used by Bhalla (2002, 2004) and Sala-i-Martin (2002, 2006), involves converting national accounts per capita income data into a common currency using PPP exchange rates, and then using information on income quintile shares to impute an income distribution for each country, for comparison against the IPL.⁷

There has been much debate regarding the appropriate method for estimating world poverty, and we do not attempt to resolve the issues raised by authors such as Atkinson and Brandolini (2001), Bhalla (2004), Chen and Ravallion (2001), Deaton (2001, 2005), Ravallion (2003a,b), Reddy and Pogge (2003) and Sala-i-Martin (2006). Rather, we focus on an aspect that is common to both approaches outlined above (and, surprisingly, has not been addressed in the global poverty debate), namely, the impact of the choice of PPP method for converting currencies. In particular, our aim is to contrast results for global poverty using our preferred EKS estimates of PPP with results obtained using the Geary-Khamis method, which has been used for global poverty measurement by Bhalla (2002, 2004),

Sala-i-Martin (2002, 2006), Berry, Bourguignon, and Morrisson (1983) and Bourguignon and Morrisson (2002).

In order to compare the impact of PPP methods, we need to choose a benchmark method for estimating global poverty. Lacking access to the expenditure survey data used by the World Bank, we decided to follow the second approach outlined above.⁸ In particular, we follow Sala-i-Martin (2006) by using income distribution data from World Bank (2003), which updates Deininger and Squire (1996), and constructing estimates of the centile income shares using a non-parametric kernel density function.⁹ We therefore implement the \$ per day measures in terms of income, rather than consumption, and we anchor our international comparisons on national accounting measures of average income rather than expenditure as measured in surveys. Our use of income poverty based on national accounts data rather than consumption poverty based on household survey data necessarily implies that our estimates of the number of people in poverty will be lower than World Bank estimates. We are aware that there is contentious debate about global poverty measurement, but our choice of method is based on the pragmatic grounds of data availability, and should not be interpreted as a rejection of arguments that favour survey-based approaches.

We have attempted to construct data that is comparable to that used by Sala-i-Martin (2006) to enhance comparability and hence where possible have used the same sources. However, there are several key differences that need to be mentioned. First, there are marked differences in the coverage of countries between our studies. While we started out with a total of 117 countries – 115 ICP countries plus China and India – we had to drop 14 of the ICP countries because we could not obtain quintile share data for them from either Deininger and Squire (1996) or the World Development Indicators - WDI - (World Bank, 2003). We also decided to drop six other ICP countries because the quintile information was either too old or considered unreliable by the source (see the Appendix for more details). The combined population of the 20 countries dropped from our data is 72 million and our final data set of 97 countries accounts for a total population of 5 billion, comprising 87 percent of the 1996 world population of 5.8 billion.¹⁰ Sala-i-Martin's(2006) dataset of 138

countries accounted for 93 percent of the world population in 2000.

Some sample selection bias may be introduced by our criteria which exclude some very poor countries for which quintile share data could not be found, whereas Sala-i-Martin (2006) introduces a different source of bias by using data from neighbouring countries to estimate these shares. But given that our different criteria account for only six percent of the world's population, we are confident that the results on which we focus – the sensitivity of poverty estimates to the choice of PPP method – are qualitatively correct if not quantitatively exact.

The second significant difference is that we do not use the PWT estimates of real GDP per capita; rather we derive our GK and EKS indexes directly from the ICP data in order to highlight the differences between the two methods. This approach has required us to produce our own estimates of real GDP for India and China, which were not included in the 1996 ICP surveys. We have imputed EKS real GDP per capita for these countries by regressing the EKS index numbers on three 1996 variables from the PWT6.1 data set: i) real GDP per capita; ii) the ratio of total trade flows to GDP ($OPEN_i$); and iii) the ratio of the price of investment goods to the price of consumption goods, $(PI/PC)_i$:¹¹

$$\ln y_i^{EKS} = -0.158 + 1.029 \ln y_i^{PWT} - 0.001 OPEN_i - 0.052 (PI/PC)_i \quad (5)$$

(0.135) (0.014) (0.000) (0.015)

$$R^2 = 0.986, \quad n = 113$$

The coefficient on real GDP captures systematic bias related to relative income levels conditional on the impact of trade and domestic price differentials. Variations in $OPEN_i$ and $(PI/PC)_i$ are indicators of the extent to which a country's price structure is likely to differ from other countries' prices, and hence the extent of substitution bias.

The PWT 1996 per capita real GDP measures for China and India are \$2969 and \$2118. Since the PWT method tends to overstate the real incomes of poorer countries, it is no surprise to find that the imputed EKS estimates are somewhat lower at \$2848 and \$2023,

respectively.

While the PWT real GDP is constructed using the GK approach, the PWT's use of super-country weights means the PWT-GK real GDP numbers are not directly comparable to the ICP-GK numbers that we have constructed here, and so we need to impute ICP-GK numbers for China and India. A regression of (log) GK real GDP on (log) PWT real GDP produces the following results:

$$\ln y_i^{GK} = 0.118 + 0.996 \ln y_i^{PWT} \quad R^2 = 0.995, \quad n = 115 \quad (4)$$

(0.086) (0.010)

The imputed GK per capita real GDP values for China and India are \$3230 and \$2308, respectively.¹² Note that the ICP-GK estimates are systematically higher than the PWT real GDP, and this has implications for the comparability of our GK poverty estimates with those of Sala-i-Martin.

5 Global poverty analysis

For our poverty analysis, we use the poverty lines of Sala-i-Martin (2006), which are updates of the \$1/day and \$2/day absolute poverty lines originally used by the World Bank: the \$1/day poverty line translates to \$532 per year in 1996 dollars, while the \$2/day line translates to \$1064 per year.¹³

A summary of our results is presented in Tables 1 and 2. It is apparent that the choice of PPP index has a marked impact on estimates of world poverty. The number of people estimated to live in extreme poverty rises by 44 percent when we switch from GK to EKS purchasing power parities. With the \$2 poverty line, the estimated number of people in poverty rises by 29 percent when we switch from GK to EKS purchasing power parities.

It is of interest to see whether the different purchasing power parity methods lead to markedly different conclusions regarding the incidence of poverty in different regions and the distribution of poverty across regions. The extent of underestimation of poverty arising from the Geary-Khamis method depends on the extent of substitution bias, which varies

Table 1: Global poverty estimates (1996) - \$1/day poverty line

	Population		<i>Geary-Khamis real GDP</i>			<i>EKS real GDP</i>		
			Poverty			Poverty		
	(mill.)	(%)	rate (%)	No. of poor (mill.) (%)		rate (%)	No. of poor (mill.) (%)	
Africa - Sub Saharan	278	5.6	36.5	101	68.0	43.0	120	55.9
East Asia & Pacific	1619	32.4	0.9	15	9.9	1.8	30	13.8
East. Europe & Central Asia	461	9.2	0.4	2	1.4	0.9	4	2.0
Latin America & Caribbean	344	6.9	1.2	4	2.9	2.4	8	3.9
Middle East & Nth. Africa	175	3.5	3.1	5	3.6	4.9	9	4.0
South Asia	1233	24.7	1.7	21	14.2	3.5	43	20.2
High-income	890	17.8	0.0	0	0.0	0.0	0	0.0
World (n=97)	5000	100.0	3.0	149	100.0	4.3	214	100.0

Table 2: Global poverty estimates (1996) - \$2/day poverty line

	Population		<i>Geary-Khamis real GDP</i>			<i>EKS real GDP</i>		
			Poverty			Poverty		
	(mill.)	(%)	rate (%)	No. of poor (mill.) (%)		rate (%)	No. of poor (mill.) (%)	
Africa - Sub Saharan	278	5.6	62.1	173	23.5	68.5	190	20.1
East Asia & Pacific	1619	32.4	13.8	223	30.4	18.2	295	31.1
East. Europe & Central Asia	461	9.2	3.0	14	1.9	5.3	25	2.6
Latin America & Caribbean	344	6.9	12.3	42	5.8	14.5	50	5.3
Middle East & Nth. Africa	175	3.5	8.0	14	1.9	11.4	20	2.1
South Asia	1233	24.7	21.8	269	36.6	30.0	369	38.9
High-income	890	17.8	0.0	0	0.0	0.0	0	0.0
World (n=97)	5000	100.0	14.7	735	100.0	19.0	949	100.0

substantially from country to country – as illustrated in Figures 2 and 3. It also depends on the shape of the income density function, particularly on ‘bunching’ of the population at income levels close to the poverty lines.

For our regional analysis, we have identified six regions that broadly coincide with the regional groupings used by the World Bank: Sub-Saharan Africa (SSA), East Asia & Pacific (EAP) which includes China, Eastern Europe & Central Asia (ECA), Latin America & Caribbean (LAC), Middle East & North Africa (MENA), and South Asia (SA) which includes India. For completeness and interpretation we also include as a separate grouping the high-income countries (HIC) from our data set. See the Appendix for a complete listing

of all countries and their regional classification.

Using the \$1/day poverty line, Africa has the majority of the world's extremely poor people and the highest regional poverty rate whether we use the GK or the EKS approach. The GK method underestimates the number of extremely poor people in Africa by 16%, yielding an estimate of 101 million compared with the more reliable EKS estimate of 120 million. In East Asia and in South Asia, however, the GK estimates understate the number of the extreme poor by more than half, with the estimates rising from 15 million and 21 million to 30 million and 43 million respectively. This indicates that there is significant bunching of the population in East and South Asia just above the GK \$1/day poverty line. The main point of this is not that we should be any less concerned about extreme poverty in Africa, but rather that estimates of the regional composition of world poverty are highly sensitive to the choice of index number approach for constructing real income.

Using the \$2/day poverty line we find that the choice of index number method has a smaller but still significant impact on regional poverty rates and the distribution of the poor. A shift from GK to EKS results in a 6.4 percentage point increase in the poverty rate in Africa compared with a 4.4 point increase in East Asia & Pacific, and a 8.2 point increase in South Asia. Under the GK approach, 67 percent of the world's poor are estimated to live in Asia, while Africa is home to 24 percent of the poor. Switching to the EKS approach, Africa's share of the poor decreases by 3.2 percentage points and South Asia's share increases by 2.3 percentage points.

6 Conclusions

The existence of substitution bias in fixed-price index numbers implies that purchasing power parity measures that rely on the Geary-Khamis method are not appropriate for measurement of global poverty. Because the GK method values incomes at prices corresponding to those found in relatively rich countries, it tends to overstate real incomes in the poorest countries relative to the US by an order of magnitude averaging twenty

percent (but up to fifty percent for some countries). On the other hand, the EKS method exhibits no systematic bias and appears to be appropriate for assessing the extent and distribution of poverty.

Our results suggest that the findings of studies which have relied on either Penn World Table data or Maddison's data to estimate levels and trends in world poverty are likely to be unreliable. Switching to the EKS method from the GK method raises the estimated number of the very poor by nearly 45 percent. We find that Asia is home to nearly 60 percent of the people who are misclassified as living just above the poverty line when the GK method is used. Given the substantial and variable degree of substitution bias that we have documented in the GK methodology, our general conclusion appears robust. The index number approach matters to global poverty measurement, and it matters a lot.

Notes

¹These indexes are based on the work of Eltetö and Köves (1964), Szulc (1964), Geary (1958) and Khamis (1972) respectively.

²See Heston, Summers, and Aten (2002) for details of the latest version, PWT6.1.

³The PWT v6.1 is based on the GK method but it uses a ‘super-country’ weighting system which is slightly different from the procedure described above in that it inflates the populations of ICP countries, classed in seven income groups, to the world population in each income group. We find that the average degree of substitution bias is unaffected by the use of super-country weights. This may be due to the fact that while India and China constitute 40% of the world’s population, they produce only 15% of world GDP.

⁴The ICP data can be accessed at <http://pwt.econ.upenn.edu/Downloads/benchmark/benchmark.html>. See Table A1 in the Appendix for a list of the 115 countries. In the Appendix – available at <http://hdl.handle.net/1885/45735> – we discuss some apparent errors in the ICP price data and we explain the assumptions we adopted to rectify these errors. Lacking access to accurate ICP data, our results must be treated with some caution, however our corrections are associated with items that comprise only a very small part of expenditure on GDP. We are primarily concerned with differences between the EKS and GK measures which we construct from the same data, so we expect our results to be robust.

⁵Since the World Bank also uses the EKS method, we compared its real GDP numbers (taken from World Bank, 2003) with our EKS index. The regression results for a reduced sample of 113 countries for 1996 (the World Bank data do not include Bermuda and Qatar) are: $\ln y_i^{GK} = 0.780 + 0.933 \ln y_i^{WB}$ $R^2 = 0.976$, $n = 113$ where $\ln y^{WB}$ is the log of the World Bank’s real GDP measure (standard errors have been

omitted for brevity). The corresponding regression using our log EKS income (with the reduced sample size) is: $\ln y_i^{GK} = 0.695 + 0.938 \ln y_i^{EKS}$ $R^2 = 0.990$, $n = 113$. In terms of the implied estimate of the degree of bias in the GK index, the World Bank's real income index is therefore very similar to our EKS index. The high degree of similarity of the two indexes is also evident from the following regression: $\ln y_i^{WB} = 0.031 + 0.991 \ln y_i^{EKS}$ $R^2 = 0.986$, $n = 113$.

⁶Calculated as $\exp[-(1 - 0.940) \times \ln(1/50)] = 1.265$.

⁷Note that Bhalla and Sala-i-Martin use differing imputation techniques and data sources to generate their estimates of the world distribution of income. Bourguignon and Morrisson (2002) also combine national accounts and survey data in their analysis of global poverty.

⁸Although grouped income and consumption data from the World Bank's surveys is public domain through the PovcalNet website, the underlying unit-record data is not.

⁹We note that it would have been possible to use the grouped income or expenditure data that are available via the World Bank's PovcalNet website, rather the income quintile data in World Bank (2003), in this exercise. However, we felt that given our main purpose is to analyse the impact of PPP methodology on global poverty measurement, and not to measure global poverty *per se*, it was important to use the same data source as Sala-i-Martin (2006), since we are using his approach for imputing income distributions. However, we note that Atkinson and Brandolini (2001) have identified problems of comparability across observations in the Deininger and Squire (1996) data, and there are other issues with using secondary datasets, which we ignore here.

¹⁰Our full data set of 117 countries covers a combined population of 5.1 billion (88 percent of world population).

¹¹Hong Kong and Singapore were excluded from the regression because their huge

trade ratios mainly reflect their entrepôt activities – warehousing imports to be shipped on to other countries.

¹²We compared our imputed estimates for India and China with the World Bank PPP estimates for 1996, extracted from the World Development Indicators. We note that the Bank estimates for these countries are imputations, but they give no information on the exact approach. Our EKS ratios to the USA are 6.9% and 9.8% respectively i.e. slightly below the World Bank’s in the case of India (7.8%) and conversely for China (9.3%). These results indicate that using the World Bank’s EKS index and imputation method would not change our results significantly (computations available from the authors by request).

¹³In contrast, the 1993-based World Bank poverty line used by Chen and Ravallion (2001, 2007) and in work measuring progress towards the Millennium Development Goals was constructed not by adjusting the 1985 poverty line for inflation, but by repeating the original process undertaken to construct the 1985 poverty line, i.e., approximating the poverty lines of the poorest countries (see Chen and Ravallion, 2001).

References

- AFRIAT, S. (1967): “The Construction of a Utility Function from Expenditure Data,” *International Economic Review*, 8(1), 67–77.
- (1984): “The True Index,” in *Demand, Equilibrium, and Trade: Essays in Honor of Ivor F Pearce*, ed. by A. Ingham. St. Martin’s Press, New York.
- ATKINSON, A., AND A. BRANDOLINI (2001): “Promise and Pitfalls in the Use Of ‘Secondary’ Data-Sets: Income Inequality in OECD Countries as a Case Study,” *Journal of Economic Literature*, 39(3), 771–799.
- BALASSA, B. (1964): “The Purchasing Power Parity Doctrine: a Reappraisal,” *Journal of Political Economy*, 72(6), 584–96.
- BERRY, A., F. BOURGUIGNON, AND C. MORRISSON (1983): “Changes in the World Distribution of Income between 1950 and 1977,” *Economic Journal*, 93(37), 331–350.
- BHALLA, S. S. (2002): “Imagine There’s No Country: Poverty, Inequality and Growth in the Era of Globalization,” Washington, DC, Institute for International Economics.
- (2004): “Poor Results and Poorer Policy: A Comparative Analysis of Estimates of Global Inequality and Poverty,” *CESifo Economic Studies*, 50(1), 85–132.
- BOURGUIGNON, F., AND C. MORRISSON (2002): “Inequality among World Citizens: 1820-1992,” *American Economic Review*, 92(4), 727–744.
- CHEN, S., AND M. RAVALLION (2001): “How Did the World’s Poor Fare in the 1990s?,” *Review of Income and Wealth*, 47(3), 283–300.
- (2007): “Absolute Poverty Measure for the Developing World, 1981-2004,” World Bank Policy Research Working Paper 4211.
- DEATON, A. (2001): “Counting the World’s Poor: Problems and Possible Solutions,” *World Bank Research Observer*, 16, 125–147.

- (2005): “Measuring Poverty in a Growing World (or Measuring Growth in a Poor World),” *Review of Economics and Statistics*, 87(1), 1–19.
- DEININGER, K., AND L. SQUIRE (1996): “A New Data Set Measuring Income Inequality,” *World Bank Economic Review*, 10(3), 565–91.
- DOWRICK, S., AND M. AKMAL (2005): “Contradictory Trends in Global Income Inequality: A Tale of Two Biases,” *Review of Income and Wealth*, 51(2), 201–29.
- DOWRICK, S., AND J. QUIGGIN (1994): “International Comparisons of Living Standards and Tastes: A Revealed Preference Approach,” *American Economic Review*, 84, 332–41.
- (1997): “True Measures of GDP and Convergence,” *American Economic Review*, 87(1), 41–64.
- ELTETÖ, O., AND P. KÖVES (1964): “On a Problem of Index Number Computation Relating to International Comparisons,” *Statisztikai Szemle*, 42, 507–518, [in Hungarian].
- GEARY, R. (1958): “A Note on the Comparison of Exchange rates and Purchasing Power Between Countries,” *Journal of the Royal Statistical Society (Series A)*, 121, 97–99.
- HESTON, A., R. SUMMERS, AND B. ATEN (2002): “Penn World Table Version 6.1,” Center for International Comparisons at the University of Pennsylvania (CICUP), October.
- HILL, R. (2000): “Measuring Substitution Bias in International Comparisons Based on Additive Purchasing Power Parity Methods,” *European Economic Review*, 44(1), 145–162.
- KHAMIS, S. (1972): “A New System of Index Numbers for National and International Purposes,” *Journal of the Royal Statistical Society (Series A)*, 135, 96–121.
- MADDISON, A. (1995): *Monitoring the World Economy: 1820-1992*. Organisation for Economic Co-operation and Development, Paris and Washington, D.C.
- (2003): *The World Economy: Historical Statistics*. Organisation for Economic Co-operation and Development, Paris and Washington, D.C.

- RAVALLION, M. (2003a): "How Not to Count the Poor? A Reply to Reddy and Pogge," mimeo. World Bank, Washington DC.
- (2003b): "Measuring Aggregate Welfare in Developing Countries: How Well Do National Accounts and Surveys Agree?," *Review of Economics and Statistics*, 85(3), 645–52.
- REDDY, S., AND T. POGGE (2003): "How *not* to Count the Poor," available at: <http://www.socialanalysis.org>.
- SALA-I-MARTIN, X. (2002): "The World Distribution of Income (estimated from Individual Country Distributions)," National Bureau of Economic Research, Inc, NBER Working Papers: 8933.
- (2006): "The World Distribution of Income: Falling Poverty and .. Convergence, Period!," *Quarterly Journal of Economics*, 121(2), 351–397.
- SAMUELSON, P. (1964): "Theoretical Notes on Trade Problems," *Review of Economics and Statistics*, 46(2), 145–54.
- SZULC, B. (1964): "Index Numbers for Multilateral Regional Comparisons," *Przegląd Statystyczny*, 3, 239–254, [in Polish].
- VARIAN, H. (1983): "Non-Parametric Tests of Consumer Behaviour," *Review of Economic Studies*, 50(1), 99–110.
- WORLD BANK (2000): *World Development Report 2000/2001: Attacking Poverty*. Oxford University press, New York.
- (2002): *Globalization, Growth and Poverty*. Oxford University press, New York.
- (2003): "World Development Indicators (WDI): 2003," CD-ROM, World Bank.

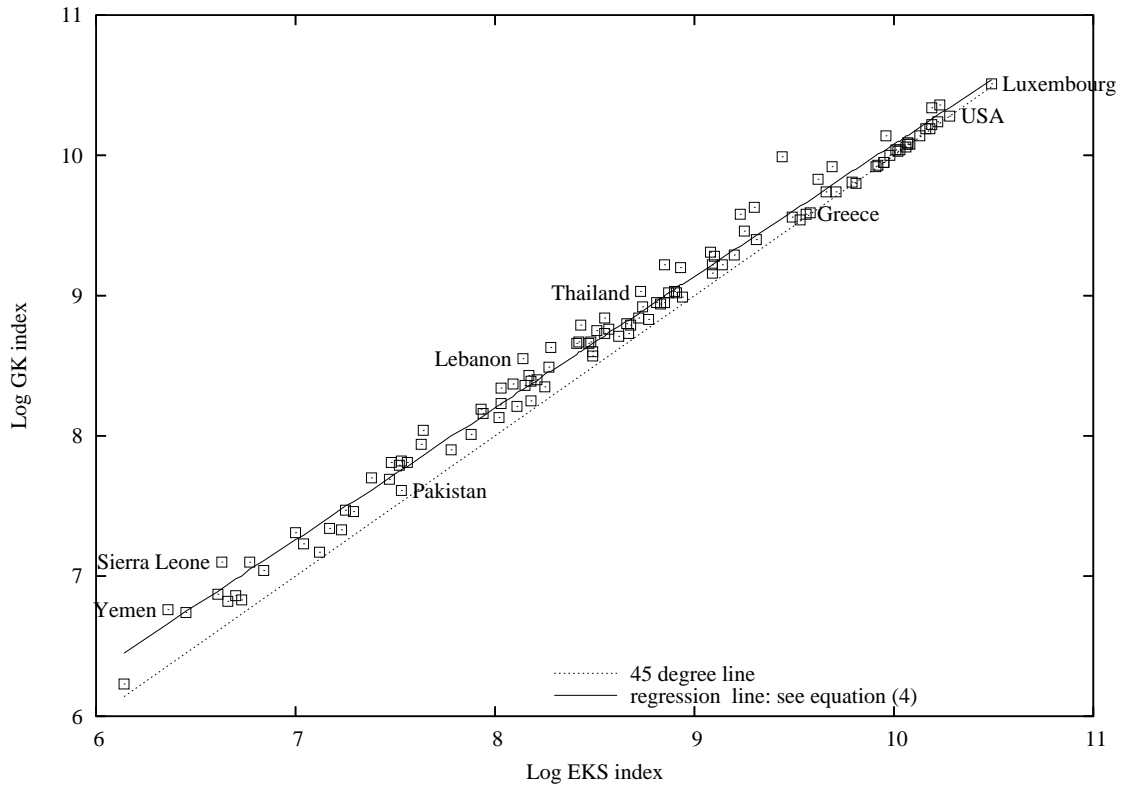


Figure 1: Comparison of the GK and EKS indexes

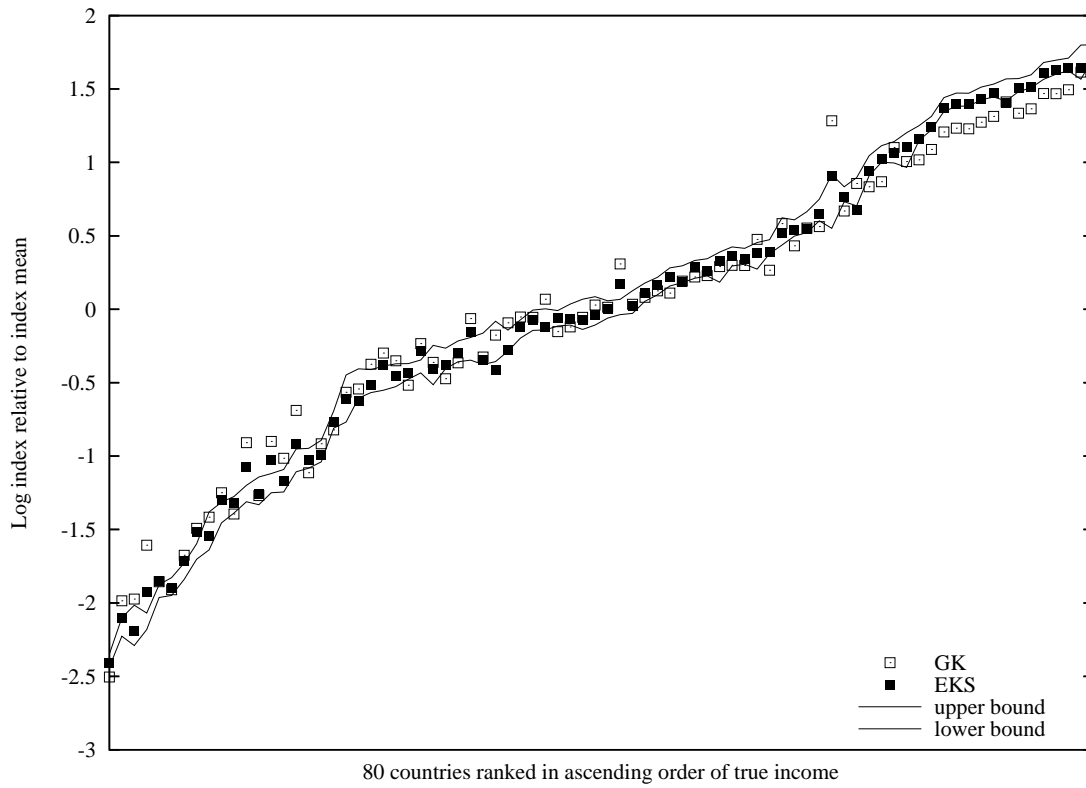


Figure 2: Afriat Bounds, GK and EKS

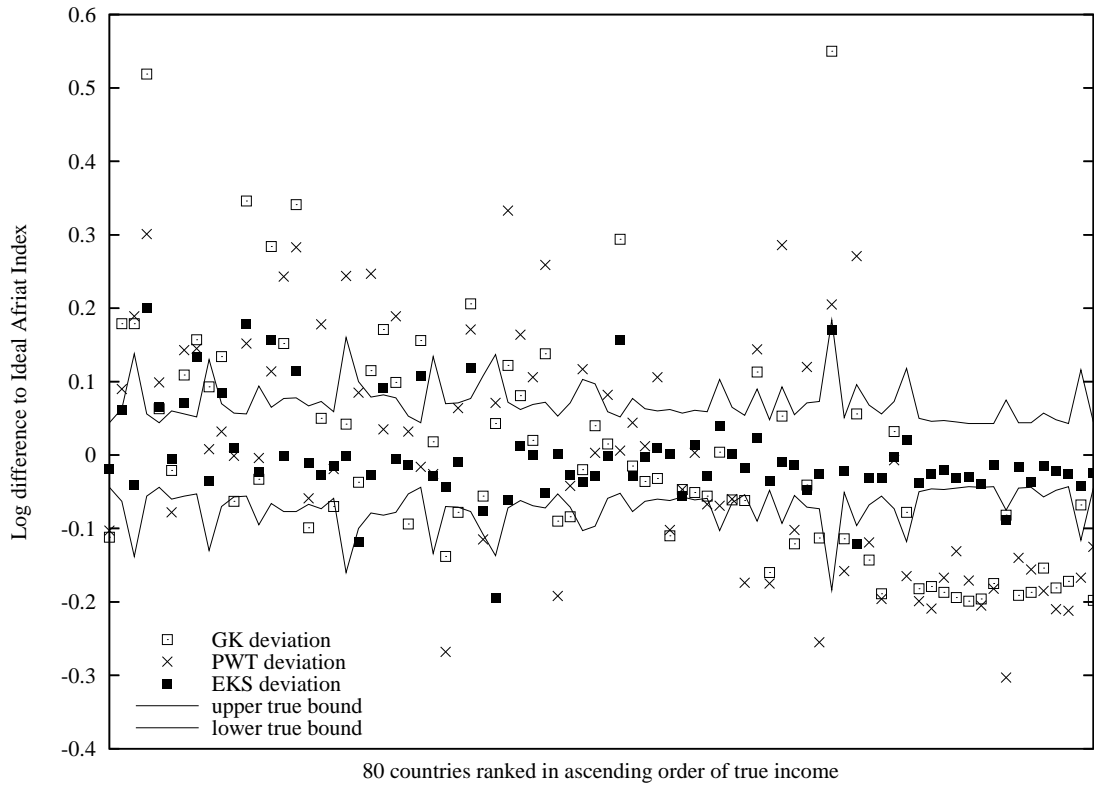


Figure 3: Deviation of PWT, GK and EKS from True Bounds