

MEASURING GLOBAL POVERTY: WHY PPP METHODS MATTER

Robert Ackland, Steve Dowrick, and Benoit Freyens*

Abstract—We present theory and evidence to suggest that, in the context of analyzing 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 and 2005 International Comparison Project data confirms that the Geary-Khamis approach substantially overstates the relative incomes of the world's poorest nations, and this leads to misleading comparisons of poverty across regions and over time. The EKS index of real income is much closer to being a true index of economic welfare and is therefore preferred for assessment of global poverty.

I. 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 by authors such as Atkinson and Brandolini (2001), Bhalla (2004), Chen and Ravallion (2001), Deaton (2001, 2005), Ravallion (2003a, 2003b), 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-Khamis (GK) index on which both the Penn World Table (Heston, Summers, & Aten, 2009) and Angus Maddison (1995, 2003) base their measures of real GDP.¹ Since 2000, the World Bank's purchasing power parity estimates used for measuring global poverty and inequality have not been based on the Geary-Khamis method, but instead are constructed with the EKS method (Chen & 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 they 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 Geary-Khamis method understates the number of the world's poor relative to the EKS method, and that the EKS method produces income measures that are closer to economic definitions of true welfare. While the direction of the GK bias is relatively well known, its impact on global estimates of poverty is largely undocumented. We fill this gap by using the last two rounds of International Comparison Program (ICP) data on worldwide purchasing power parities (1996, 2005) to empirically quantify the impact of using different index number methods on the measurement of global poverty.

It is important to note that we do not attempt to provide a definitive count of the world's poor or offer judgments 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 II, we explain the difference between the Geary-Khamis and the EKS methods of calculating purchasing power parities. We show why the Geary-Khamis method may exaggerate the purchasing power of poor country currencies and discuss the properties of the EKS index. 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.

In section III, we apply the two purchasing power parity methodologies to 1996 and 2005 International Comparison Project data. Our analysis confirms the findings of Hill (2000), who reports that the Geary-Khamis method imparts biases of 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 purchasing power parity methods, we then demonstrate the implications of this bias for global poverty estimates. In section IV, we discuss alternative approaches to measuring world poverty and explain our choice, which involves estimating the cumulative income distribution function for a subset of the International Comparison Project countries using published data on quintile shares. We then demonstrate the impact of changing from Geary-Khamis to EKS measures of purchasing power parity in terms of estimating the

Received for publication September 19, 2008. Revision accepted for publication February 2, 2012.

* Ackland: Australian Demographic and Social Research Institute, Australian National University; Dowrick: School of Economics, Australian National University; Freyens: Faculty of Business, Government, and Law, University of Canberra and Centre of Law and Economics, Australian National University.

We thank two anonymous referees for providing helpful comments on an earlier version of this paper.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00294.

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

total number of people in poverty, the distribution of the poor across and within regions, and temporal changes in poverty.

II. 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, referred to as the Balassa-Samuelson effect (Balassa, 1964; Samuelson 1964). Real wages are low in countries with low labor productivity, so nontraded labor-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), which provides measures of real GDP per capita at constant international prices for over one hundred countries since 1950.² These data are commonly referred to as PPP measures of real income.

The PWT estimates of average real income (GDP per capita) are based on the Geary-Khamis (GK) method of constructing average international prices, with the GDP of each country being evaluated at these fixed prices. Formally, GK measures of real GDP and the corresponding purchasing power parities are derived as the solution to the following set of simultaneous equations:

$$\pi_k = \sum_i \frac{p_k^i q_k^i}{e^i \sum_i q_k^i} \quad k = 1 \dots K \text{ items of expenditure,} \quad (1)$$

$$e^i = \frac{\mathbf{p}^i \cdot \mathbf{q}^i}{\boldsymbol{\pi} \cdot \mathbf{q}^i} \quad i = 1 \dots N \text{ countries.}$$

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, comprising 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, $\mathbf{p}^i \cdot \mathbf{q}^i / e^i$.

An alternative way of comparing real incomes across countries is to use the EKS quantity index.³ This is based on the bilateral Fisher quantity index, F^{ij} , which is constructed as

the geometric mean of the Paasche quantity index (Q_P^{ij}) and Laspeyres quantity index (Q_L^{ij}):

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

The EKS quantity index (hereafter, the EKS index) is constructed as the geometric mean of Fisher quantity indexes:⁴

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

A. Evaluating Bias in the GK and EKS Indexes

The GK method succeeds in removing the traded-sector bias in exchange rate valuations, yielding real income differentials that are substantially smaller than the income differentials of exchange rate comparisons. However, because international prices are calculated 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, the GK 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.⁵ The use of a single international price vector in GDP valuations therefore introduces systematic bias by ignoring substitution toward higher consumption of goods and services that are locally cheap. Evaluated at nonlocal 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. This is referred to as the Gershenkron effect, after Gershenkron (1951).

With regard to the EKS index, substitution bias implies that the Paasche index is likely to understate the true income ratio of the richer country relative to the poorer country, while the Laspeyres index tends to overstate it. However, it is not necessarily the case that the two biases will exactly cancel out in the Fisher index (and hence the EKS index).

The existence and extent of bias in the GK and EKS indexes can be assessed using the theory of consumer behavior that 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, and this can be tested using the generalized

⁴ An alternative approach for constructing EKS real income is to deflate local currency GDP using the EKS price index (constructed analogously to the EKS quantity index presented here, but using Paasche and Laspeyres price indexes) as the PPP exchange rate (see, for example, Deaton & Heston (2010)).

⁵ In 1996, 45% of world GDP was produced by just seven OECD countries: the United States, Japan, Germany, France, United Kingdom, Italy, and Canada.

² See Heston, Summers, and Aten (2009) for details of the latest version, PWT6.3.

³ The OECD currently uses the EKS method for comparisons of incomes of their member countries.

axiom of revealed preference (GARP) (Afriat, 1967; Diewert, 1973; Varian, 1982). If we cannot reject the hypothesis that a given set of demand data was generated by a representative household, 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.⁶

Afriat (1984) has shown that the test for common homothetic preferences requires that there exists 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 .⁷ Varian (1983) provides an efficient test of common homothetic preferences (the homothetic axiom of revealed preference, HARP) that involves using Warshall's algorithm to construct the minimum path matrix using the matrix of bilateral Laspeyres indexes as input. However satisfying, the test of common homothetic preferences does not imply a unique Afriat index. Rather, there are well-defined bounds within which there is an irreducible indeterminacy resulting from the fact that we do not observe utility. Dowrick and Quiggin (1997) show that the elements of the minimum path matrix can be used to construct \mathbf{a}^- and \mathbf{a}^+ , which represent the lower and upper bounds of true indexes relative to the geometric mean (which they denote the ideal Afriat index), thus allowing a definition of bias as the extent to which the GK or EKS indexes violate the true bounds.

We use the bounds of the Afriat true index as a benchmark against which to measure bias in other indexes because it is, by construction, a utility-based measure 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 prices of each country. The representative consumer is, of course, a convenient fiction; individual consumers may well have heterogeneous and nonhomothetic preferences. We are, however, able to exploit the remarkable finding that aggregate (per capita) behavior 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.⁸

⁶ Ackland (2008) extends Afriat's work by providing tight bounds to the Allen marginal index, which measures the utility derived from above-subsistence consumption and can be constructed for demand data that are consistent with common affine-homothetic preferences.

⁷ A multilateral index satisfies circularity: the real income of country i relative to country j is the same whether the two are compared directly or via an intermediate third country k .

⁸ It has been noted by authors such as Bronars (1987) that revealed preference tests of common nonhomothetic preferences using aggregate data

III. Comparing GK, EKS, and True Income Levels

We used 1996 and 2005 ICP data to calculate three real income measures: GK, EKS, and the Afriat True index. The 1996 ICP data cover 115 countries, while the 2005 data collection extended to 146 countries.⁹ Significantly for this paper, with its focus on global poverty measurement, the 1996 data do not include India and China, and we detail below the approach we used to impute income for these populous and relatively poor countries. The 2005 ICP was the first benchmark exercise that includes both India and China.

The ICP data consist of prices and per capita expenditures in local currency on various components of GDP, based on the familiar national accounting identity: $Y = C + I + G + NX$. The 1996 data have 31 categories of goods and services, and the 2005 data, cover 37 items. From these data, we calculate the real per capita quantities of each item in each of the 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 subsample of countries, as explained later.

We construct the GK index by solving the system of equations (1), normalizing the parities so that $e^{USA} = 1$, giving a vector of national purchasing power parities, \mathbf{E}^{GK} , and a vector of real GDP per capita, \mathbf{y}^{GK} , where the value of real GDP for the United States equals its local currency value. The EKS index is constructed from equation (3), and we normalize this to the index \mathbf{y}^{EKS} , where the value for the United States equals its local currency GDP.

The two measures of log real GDP are highly correlated, as shown in the scatter plots in figure 1a (1996) and figure 1b (2005), where log GK income is on the vertical axis and log EKS income on the horizontal axis. The OLS regressions are reported in equation (4a) (1996) and equation (4b) (2005), with the standard error of coefficient estimates in brackets. The regression lines and 45 degree lines are also shown in figure 1.

$$\ln y_i^{GK} = 0.680 + 0.940 \ln y_i^{EKS} \quad R^2 = 0.990, \quad (4a)$$

(0.076) (0.009)

$$s.e. = 0.103, \quad n = 115,$$

$$\ln y_i^{GK} = 0.358 + 0.967 \ln y_i^{EKS} \quad R^2 = 0.996, \quad (4b)$$

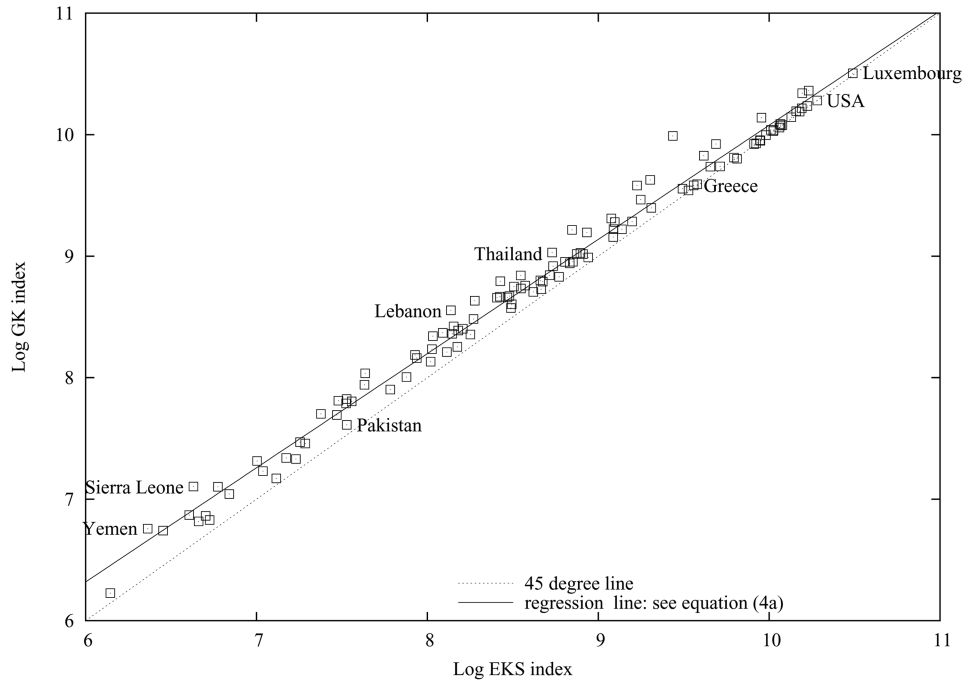
(0.044) (0.005)

$$s.e. = 0.082, \quad n = 147.$$

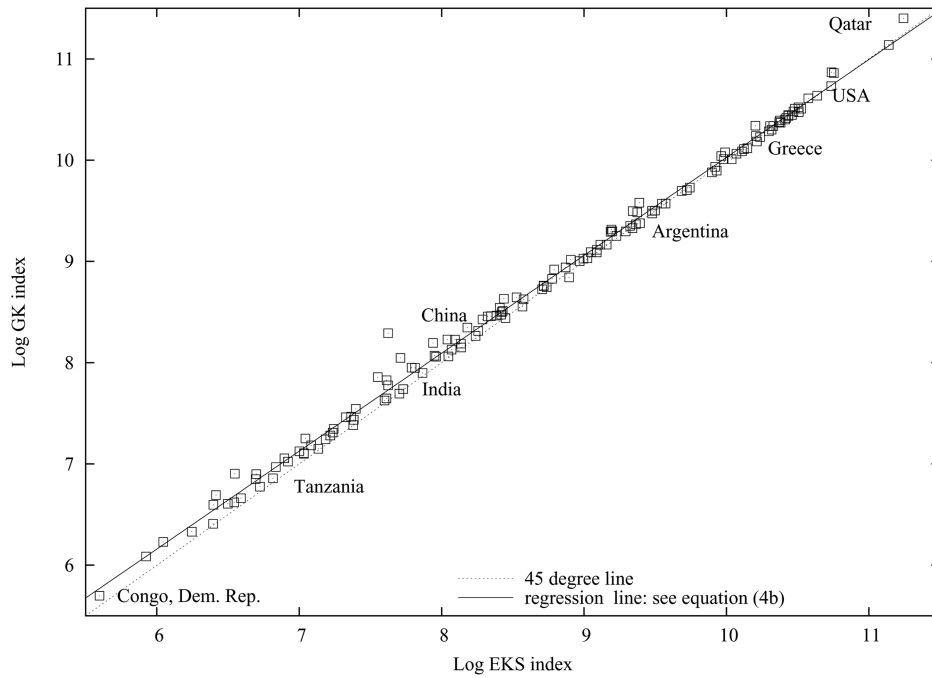
may lack power because large disparities in income and smaller differences in relative prices mean that budget lines rarely intersect. To address this, Blundell, Browning, and Crawford (2003) use nonparametric Engel curves to estimate expansion paths, allowing the projection of consumer preferences onto new budget lines that are parallel shifts of the actual budget lines; with more intersections of budgets, the power of the GARP test of common nonhomothetic preferences can be increased. However, the HARP test of common homothetic preferences does not suffer from low power caused by rarely intersecting budgets since under homotheticity, expansion paths are rays through the origin, and the projection of consumer preferences onto parallel budget lines is implicit in the test.

⁹ The supplemental appendix provides details on the data sources and preparation.

FIGURE 1.—COMPARISON OF THE GK AND EKS INDEXES



(a) 1996



(b) 2005

For the 1996 data, 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. For a poor

country such as Yemen, with real income one-fiftieth that of the United States, the 1996 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 1.¹⁰ The regression slope coefficient for 2005 (0.967) is also significantly less than unity, but the fact that it is larger than the 1996 slope coefficient indicates that the 2005 GK index is not compressing the cross-country income distribution as much relative to the EKS index (and this is visually apparent in figure 1).

We can use the bounds to the Afriat True index to assess the magnitude of bias in the GK and EKS indexes within the set of countries for which aggregate homotheticity is not rejected. Following Varian (1983), we used Warshall's algorithm to construct the minimum path matrix for various sets of countries, with the hypothesis of common homothetic preferences being rejected if any elements of the diagonal were negative. We used an iterative procedure to find the maximal set of countries satisfying the test of common homothetic preferences, and for the 1996 data, we identified a set of eighty countries for which we cannot reject the hypothesis of common homothetic preferences. 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. Around the same proportion of countries in the 2005 ICP data satisfied the test; our iterative approach identified a set of 104 countries not rejecting aggregate homotheticity.

The bias in the GK and EKS indexes is graphically presented in figure 2a (1996) and figure 2b (2005), which show the two real income measures together with the upper and lower Afriat bounds (\mathbf{a}^- and \mathbf{a}^+), where all are expressed as deviations from the midpoint of the Afriat bounds. The countries are ordered from the poorest to the richest along the horizontal axis.

For 1996, the GK index violates the bounds for 53 cases out of 80, and the mean absolute deviation outside the bounds is 17%. The GK index provides a very obvious overestimation of real GDP for the poorest countries and underestimation for the richest. When we examine the deviations of the EKS measure, we find that 80% of the observations lie within the true bounds. The 16 deviations outside the bounds are smaller, averaging 12%. Furthermore, there is no tendency for the EKS to systematically under- or overestimate the true GDP of the relatively rich or relatively poor.

The 2005 GK index violates the bounds 84 times (a greater proportion of violations than found with the 1996 data), but the mean bias is less, at 10%. The EKS index violates the bounds 40 times (almost twice the proportion of violations found in 1996), but the mean absolute deviation of 4% is considerably smaller than that found for the GK index (and it is also smaller than the bias for the EKS index in 1996). Although the GK is less biased in 2005 (compared with what was found for 1996), there is still clear visual evidence that this index systematically underestimates the income of rich countries and overestimates that of poor countries, while the EKS index does not exhibit systematic bias.

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 the GK index for our subsample of countries satisfying homotheticity, we regard the EKS index as far more suitable for the measurement of world poverty.¹¹ The GK income measures are likely to substantially overstate true income levels in the world's poorest countries relative to the United States, implying that any poverty measurement based on the GK is likely to substantially understate the extent of poverty when the international poverty line is defined in U.S. dollars. Moreover, since the degree of bias in GK income varies from country to country, poverty estimates based on the GK are likely to misallocate the world's poor across countries.

As a final point, it should be noted that while we have focused on the standard construction of the GK index in the analysis, the PWT uses a variant of the GK that involves a "supercountry" weighting system that inflates the populations of ICP countries, classed in a number of income groups, to the world population in each income group. However, we found that for the 1996 data, the average degree of substitution bias is unaffected by the use of supercountry weights. This may be due to the fact that while India and China constituted 40% of the world's population in 1996, they produced only 15% of world GDP. Below, we provide further discussion relating to the situation in 2005.

A. Imputing Real Income for China and India in 1996

The 1996 ICP did not include China or India, omitting over 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 the ICP sample and other surveys. While the PWT real GDP is a variant of the GK index, we did not use it in our global poverty analysis, preferring to derive our GK and EKS indexes directly from the ICP data in order to highlight the differences between the two methods. We therefore used a two-step procedure involving PWT data to estimate the GK and EKS measures for both China and India, giving us our extended ICP sample of 117 countries, which includes 91% of the world's population and 95% of the world's GDP in 1996.


The first step was to impute EKS real GDP per capita by regressing the EKS index numbers on three variables from the 1996 PWT6.1 data set: (a) real GDP per capita (y_i^{PWT}), (b) the ratio of total trade flows to GDP ($OPEN_i$), and (c) the ratio of the price of investment goods to the price of consumption goods, $(PI/PC)_i$:¹²

¹¹ Neary (2004) questions the use of the EKS index in multilateral comparisons when preferences are not homothetic and proposes a modification of the GK index (the Geary-Allen international accounts system) that minimizes substitution bias. While noting these theoretical and empirical contributions, we focus on the EKS index here because of its prominent role in the international comparisons work of the World Bank and other international agencies.


¹² Hong 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.

¹⁰ Calculated as $\exp[-(1 - 0.940) \times \ln(1/50)] = 1.265$.

FIGURE 2.—DEVIATION OF GK AND EKS FROM TRUE BOUNDS



(a) 1996



(b) 2005

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

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

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

(5) The coefficient on y_i^{PWT} captures systematic substitution 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 \$2,969 and \$2,118, respectively. Since the PWT method (as a variant of GK) tends to overstate the real incomes of poorer countries, it is no surprise to find that the imputed EKS estimates are somewhat lower, at \$2,848 and \$2,023, respectively.

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

$$\ln y_i^{GK} = 0.118 + 0.996 \ln y_i^{PWT} \quad R^2 = 0.989, \quad n = 115. \\ (0.086) \quad (0.010) \quad (6)$$

The imputed GK *per capita* real GDP values for China and India are \$3,230 and \$2,308, respectively.

Finally, we performed a quasi-validation of our imputation approaches by comparing the actual EKS and GK for each of the countries included in regressions (5) and (6) with their predicted real income levels, using the estimated coefficients from the regressions of the other $n - 1$ countries. While the average percentage prediction error was the same for both EKS and GK (1%), the GK prediction was overall marginally better, with a standard deviation of the prediction error of 11% (compared with 14% for EKS). We further compared our imputed 1996 EKS indexes for China and India with World Bank PPP estimates extracted from the World Development Indicators.¹³ Our EKS ratios to the United States are 6.9% and 9.8%, respectively—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.

IV. Global Poverty Measurement

A. The International Poverty Line

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 a day per person" (extreme poverty) and "\$2 a day per person" (moderate poverty) IPLs (see, for example, World Bank, 2000) provide the yardstick by which most global poverty comparisons are produced today.

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

day line translates to \$1,064 per year.¹⁴ For our poverty analysis using the 2005 ICP data, we follow the same approach using recent \$US CPI to update the \$1 per day and \$2 per day absolute poverty lines to 2005 dollars: the \$1 per day poverty line now translates to \$662.50, per year in 2005 dollars, while the \$2 per day line translates to \$1,325 per year.

B. Approaches for Estimating World Poverty

Two main approaches 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 & 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.¹⁵

There has been much debate regarding the appropriate method for estimating world poverty,¹⁶ 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, 2003b), 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): 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 GK method, which has been used for global poverty measurement by Bhalla (2002, 2004), Sala-i-Martin (2002, 2006), Berry et al. (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. We decided to follow Sala-i-Martin (2006) by using income distribution data to construct estimates of the centile income shares using a nonparametric kernel density function. However, we note that instead of the Sala-i-Martin

¹⁴In contrast, the 1993-based World Bank poverty line used by Chen and Ravallion (2001, 2007) and in work measuring progress toward 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 that is, approximating the poverty lines of the poorest countries (see Chen & Ravallion, 2001).

¹⁵Note 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.

¹⁶As noted by a referee, the national accounts method adopted by Sala-i-Martin is almost certain to have a large downward bias relative to any given poverty line, so that switching from this method to the survey-based approach described in Chen and Ravallion (2001) is somewhat equivalent to raising the international poverty line in the Sala-i-Martin approach.

¹³We note that the World Bank estimates for these countries are imputations, but they give no information on the exact approach.

approach, we could have used the World Bank survey-based poverty estimation approach available via the bank's PovcalNet website.¹⁷ We are aware that there is contentious debate about global poverty measurement, and our choice of method should not be interpreted as a rejection of arguments that favor survey-based approaches.

For 1996, we used income distribution data from the World Development Indicators (WDI, World Bank, 2003) and from Deininger and Squire (1996). For the 2005 exercise, we use data from the World Income Inequality Database (WIID) of the World Institute for Development Economics Research, UNU-WIDER (2007), and from the World Bank (2009). We note that it would have been possible to use the grouped income or expenditure data that are available of PovcalNet rather than the income quantile data from the WIID and WDI databases in this exercise. However, we felt that given our main purpose—to analyze the impact of PPP methodology on global poverty measurement, and not to measure global poverty per se—it was important to use the same type of data source as Sala-i-Martin (2006), since we are using his approach for imputing income distributions.

We therefore implement the dollars-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.

While we have attempted to construct data that are comparable to those used by Sala-i-Martin (2006) for comparability, several key differences need to be mentioned. First, there are marked differences in the coverage of countries between our studies. For the 1996 exercise, we started out with 117 countries (115 ICP countries, plus China and India) but had to drop 14 of the 1996 ICP countries because we could not obtain quintile share data for them from either Deininger and Squire (1996) or the WDI. Similarly, for the 2005 exercise, we started with 146 countries (inclusive of China and India) but had to drop 16 countries because decile or quintile data were unavailable from the WIID and WDI databases. We also decided to drop six other countries common to the 1996 and 2005 ICP rounds because the quantile information was either too old or considered unreliable by the source (see the online appendix for more details).

For 1996, 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% of the 1996 world population of 5.8 billion. Our full data set of 117 countries covers a combined population of 5.1 billion (88% of world population). Sala-i-Martin's (2006) data set of 138 countries accounted for 93% of the world population in 2000. For 2005, the combined population of the 22 countries dropped from our data is 135

million, and our final data set of 124 countries accounts for a total population of 5.99 billion, comprising 93% of the 2005 world's population of 6.46 billion. The full ICP 2005 data set of 146 countries covers a combined population of 6.1 billion (95% of world population). Sala-i-Martin's (2006) data set of 138 countries accounted for 93% of the world population in 2000.

Some sample selection bias may be introduced by our criteria, which exclude some very poor countries for which quantile share data could not be found, whereas Sala-i-Martin (2006) introduces a different source of bias by using data from neighboring countries to estimate these shares. But given that our different criteria account for respectively, only 6% (1996) and 2% (2005) 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.

C. Level and Composition of Global Poverty

Our global poverty estimates are presented in table 1, with the 1996 estimates in the top half and the 2005 estimates in the bottom half. It is apparent that the choice of PPP index has a marked impact on the 1996 estimates of world poverty. The number of people estimated to live in extreme poverty rises by 44% when we switch from GK to EKS purchasing power parities, whereas moderate poverty increases by 29%. In absolute numbers, this implies that the GK method underestimates extreme world poverty by 65 million people and moderate world poverty by 214 million people. The choice of index number makes less of a difference to estimates of global poverty in 2005: the EKS estimate of the number of people living in extreme poverty is 9% higher than the GK-based estimate, while the level of moderate poverty is 7% higher. This translates into increments of 33 million and 85 million poor people, respectively.

The fact that the choice of an index number approach has a greater impact on global poverty estimates for 1996 compared with what is found for 2005 has implications for accurate measurement of global poverty trends. While it is readily apparent that global poverty increased significantly between the 1996 and 2005 ICP rounds, the increase is more marked with the GK-based poverty estimates. The global rate of extreme poverty, as derived using GK real income, doubled between 1996 and 2005 (increasing by 2.9 percentage points), while it increased only by half under the EKS method. Similarly, the GK-based moderate poverty rate increases by one-third between 1996 and 2005 (an increase of 5.1 percentage points), whereas the EKS-based moderate poverty rate increases only by about 10% (2.2 percentage points).

It is of interest to see whether the different PPP 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 GK method depends on the extent of substitution bias, which varies substantially from country to

¹⁷ <http://go.worldbank.org/7X6J3S7K90>.

TABLE 1.—GLOBAL POVERTY ESTIMATES, 1996 AND 2005

	\$1/Day Poverty Line						\$2/Day Poverty Line					
	GK Real GDP			EKS Real GDP			GK Real GDP			EKS Real GDP		
	Population (million)	Rate (%)	Number (million)	Rate (%)	Number (million)	(%)	Rate (%)	Number (million)	Rate (%)	Number (million)	(%)	
1996 estimates												
Sub-Saharan Africa	278	5.6	101	68.0	120	55.9	62.1	173	23.5	190	68.5	20.1
East Asia and Pacific	1,619	32.4	15	9.9	30	13.8	13.8	223	30.4	295	18.2	31.1
East. Europe and Central Asia	461	9.2	2	1.4	4	0.9	3.0	14	1.9	25	5.3	2.6
Latin America and Caribbean	344	6.9	4	2.9	8	2.4	12.3	42	5.8	50	14.5	5.3
Middle East and North Africa	175	3.5	5	3.6	9	4.9	8.0	14	1.9	20	11.4	2.1
South Asia	1,233	24.7	21	14.2	43	3.5	21.8	269	36.6	369	30.0	38.9
High income	890	17.8	0	0.0	0	0.0	0.0	0	0.0	0	0.0	0.0
World (<i>n</i> = 97)	5,000	100.0	149	100.0	214	4.3	14.7	735	100.0	949	19.0	100.0
2005 estimates												
Sub-Saharan Africa	664	11.1	265	74.2	284	42.7	64.4	428	36.1	442	66.5	34.8
East Asia and Pacific	1,801	30.1	13	3.6	19	1.1	11.6	209	17.6	261	14.5	20.6
East. Europe and Central Asia	444	7.4	0	0.0	1	0.2	1.0	4	0.4	8	1.9	0.7
Latin America and Caribbean	470	7.8	4	1.2	5	1.1	8.8	41	3.5	45	9.6	3.6
Middle East and North Africa	205	3.4	1	0.2	1	0.3	5.5	11	1.0	13	6.5	1.1
South Asia	1,438	24.0	74	20.6	80	5.5	34.1	490	41.4	500	34.8	39.4
High income	971	16.2	0	0.0	0	0.0	0.0	0	0.0	0	0.0	0.0
World (<i>n</i> = 124)	5,993	100.0	357	100.0	390	6.5	19.8	1185	100.0	1270	21.2	100.0

These estimates are based on national accounts data rather than household survey data, and hence are not directly comparable to World Bank global poverty estimates.

country, as illustrated in figure 2. 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 (hereafter Africa); East Asia and Pacific, which includes China; Eastern Europe and Central Asia; Latin America and Caribbean; Middle East and North Africa; and South Asia, which includes India. For completeness and interpretation, we also include as a separate grouping the high-income countries from our data set.¹⁸

Using the \$1 per day poverty line, Africa had the majority of the world's extremely poor people in 1996 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 for the extremely poor population of 101 million compared with the more reliable EKS estimate of 120 million. However, in the other regions (other than high income), GK estimates understate the number of the extremely poor by about half relative to the EKS method.

For East Asia and South Asia in particular, home to the bulk of the world's non-African poor, switching to the EKS method increases the estimate of the number of people living in extreme poverty from 15 million and 21 million to 30 million and 43 million, respectively. Accordingly, East and South Asia's share of the extremely poor increases from 24% to 34% (with Africa's share declining from 68% to 56%). The marked sensitivity of our poverty estimates for East and South Asia to the choice of index number method is due to a significant bunching of the population in these regions just above the GK \$1 per day poverty line. The main point of these observations is not that we should be any less concerned about extreme poverty in Africa (as the switch to EKS-based estimates may suggest), but rather that the 1996 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 per 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 and Pacific and a 8.2 point increase in South Asia (the latter two corresponding to increases of about one-third in the respective regional poverty rates). Under the GK approach, 67% of the world's moderate poor are estimated to live in Asia, while Africa is home to 24% 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.

Turning to the 2005 results, with the \$1 per day poverty line, Africa remains the world's poorest region regardless of

index number methods. The GK method underestimates the number of extremely poor people in Africa by 19 million, or 7% relative to the EKS method. While this underestimation is the same in absolute terms to what was found for 1996, the relative difference is smaller because the overall number of people living in extreme poverty in Africa grew between the two years. There is still substantial GK bias for East Asia, Eastern Europe and Central Asia, and for North Africa and Middle East in 2005 (with the GK-based measure of income significantly reducing regional poverty rate compared to the EKS method) but only minor bias for South Asia, where the GK-based extreme poverty rate is now a mere 10% (0.4 percentage points) lower than what is found with the EKS measure. We observe similar regional poverty patterns for moderate poverty in 2005.

Finally, as was found with estimates of the trend of overall global poverty, the choice of index number method makes a qualitative difference in our understanding of how regions are progressing in the fight against poverty. Most obvious is the fact that while the GK-based estimates show an increase in the extreme poverty rate for Africa between 1996 and 2005 (from 36.5% to 39.8%), the EKS method suggests that African extreme poverty marginally declined (from 43% to 42.7%). With the GK approach the number of extreme poor in East Asia decreased from 15 million in 1996 to 13 million in 2005 (a 13% decline), while the EKS method suggests this region did much better over the period, with a decrease of 11 million (or nearly 40%) in the number of extremely poor. The progress of South Asia in the fight against poverty is also better (or less worse) when the EKS method is used, with the number of extremely poor increasing by 53 million (around 250%) under the GK method, compared with an increase of only 37 million (86%) with the EKS method.

D. Do PPP Methods Matter for Global Poverty Measurement?

The main conclusion that can be drawn from the preceding section is that while the choice of PPP method exerts considerable influence on 1996 poverty estimates, the effect is much less prominent for 2005. This observation applies as much to global poverty rates as to regional patterns, with 2005 index number effects subdued for most regions and particularly for South Asia, relative to 1996.

Given the importance of China and India to measurement of global poverty, the narrowing of the gap between GK- and EKS-based measures of poverty may be attributable to the different treatment of these countries over the two ICP rounds. As documented above, India and China were not part of the 1996 ICP benchmark, and researchers from the PWT to ourselves have had to resort to imputation techniques for these two countries. One possibility is that our method for imputing 1996 real income measures for China and India was not sufficiently accurate.

A more plausible explanation for the narrowing of the difference between GK and EKS measures of real income (and

¹⁸ See the appendix for a complete listing of all countries and their regional classification.

hence the respective global poverty estimates) from 1996 to 2005 is that there was a decline in the Gershengkron effect, in particular for the large poorer countries. We investigated the trend between 1996 and 2005 in bilateral price similarity across the countries in our study, using the cosine of the angle between a pair of price vectors (Kravis, Heston, & Summers, 1982). For each country i , the price similarity in relation to the GK international price vector $\hat{\pi}$, estimated from equation (1), is defined as

$$S^i = \frac{\sum_{k=1}^K w_k P_k^i \hat{\pi}_k}{\sqrt{\sum_{k=1}^K w_k (P_k^i)^2 \cdot \sum_{k=1}^K w_k (\hat{\pi}_k)^2}},$$

where w_k is the share of total expenditure by all countries on category k , and S^i ranges from 0 (perfect dissimilarity) to 1 (perfect similarity). Average S^i increased from 0.86 in 1996 to 0.91 in 2005, indicating that in 2005, price structures were significantly more similar to the GK international price vector for that year compared to what was found in 1996.¹⁹

While it is beyond the scope of this paper to identify the exact reasons for an increase in price similarity (and, hence, decline in the Gershengkron effect) between 1996 and 2005, it is clear that the substantial methodological changes between the two ICP rounds would be important. As documented by the World Bank (Chen & Ravallion, 2010; Ravallion, 2008, 2010; World Bank 2008), the 2005 ICP round puts greater emphasis on the comparability of goods and services per region, with stricter quality measures, clearer commodity descriptions, productivity adjustments for the government sector, and much improved country coverage than in previous editions of the ICP.

Finally, it should be noted that the inclusion of China and India in the 2005 ICP means that the GK international price vector more closely reflects the price structures of lower- and middle-income countries. Indeed, Deaton and Heston (2009) find no pattern in substitution bias (in terms of consistent inequality relationships between the GK and EKS indexes) in 2005 for the six large countries that have substantial poverty: China, India, Indonesia, Brazil, Nigeria, and Russia. As the authors note, "The contributions of these countries to world production are now large enough to remove any consistent Gershengkron effect for countries in the middle of the income distribution" (p. 13).

In the context of these substantial methodological changes to the ICP, it is important to stress once again that it is not the aim of our paper to provide definitive estimates of global poverty rates and levels and their changes over time, but to highlight the impact of using different index number methods in generating these estimates. From this perspective, it is interesting to observe that GK-based poverty rates are more sensitive to revisions of the benchmark ICP data than

the EKS method, which produces comparatively more stable estimates of global poverty. This relative sensitivity of the GK index to the methods underlying the ICP data collection and construction results in markedly different conclusions regarding progress in reducing world poverty. In particular, the GK approach produces poverty estimates that suggest significantly worse progress in poverty reduction at the global level and within the key regions of Africa, East Asia, and South Asia.

V. Conclusion

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 United States by an order of magnitude averaging 20% (but up to 50% 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 that 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. We found that for the 1996 ICP data, the choice of index number has a significant impact on poverty estimates, with the GK method underestimating the overall number of the very poor by nearly 45% and Asia is home to nearly 60% of the people who are misclassified as living just above the poverty line when the GK method is used. The 2005 ICP data reveal a less marked difference in the poverty estimates between the GK and EKS index number approaches, and this is in line with the findings of authors such as Deaton and Heston (2009) that changes in ICP methodology have reduced the Gershengkron effect. However the fact that GK-based poverty estimates are evidently more sensitive to changes in ICP methodology leads to GK and EKS methods producing markedly different conclusions regarding the trend of global and regional poverty between 1996 and 2005, with the former approach suggesting significantly worse progress in poverty reduction between these years.

Given the substantial and variable degree of substitution bias that we have documented in the GK methodology and the changes in the level of substitution bias over time due to revisions in the ICP approach, our general conclusion appears robust. The index number approach matters to global poverty measurement, and it matters a lot.

REFERENCES

- ¹⁹ In support of our claim that this is a marked increase in price similarity, we note that Dowrick and Akmal (2005) argued that a decrease in the price similarity index from 0.954 in 1980 to 0.940 in 1991 (a change of 0.014) was evidence that "price structures became markedly less similar/more dissimilar [over the period]" (p. 213).
- Ackland, R., "Incorporating Minimum Subsistence Consumption into International Comparisons of Real Income," this REVIEW 90 (2008), 702-712.
- Afriat, S. N., "The Construction of a Utility Function from Expenditure Data," *International Economic Review* 8:1 (1967), 67-77.

- “The True Index,” in A. Ingham (Ed.), *Demand, Equilibrium, and Trade: Essays in Honor of Ivor F Pearce* (New York: St. Martin’s Press, 1984).
- Atkinson, A. B., and A. Brandolini, “Promise and Pitfalls in the Use of ‘Secondary’ Data-Sets: Income Inequality in OECD Countries as a Case Study,” *Journal of Economic Literature* 39 (2001), 771–799.
- Balassa, B., “The Purchasing Power Parity Doctrine: A Reappraisal,” *Journal of Political Economy* 72 (1964), 584–596.
- Berry, A., F. Bourguignon, and C. Morrisson, “Changes in the World Distribution of Income between 1950 and 1977,” *Economic Journal* 93 (1983), 331–350.
- Bhalla, S. S., “Imagine There’s No Country: Poverty, Inequality and Growth in the Era of Globalization” (Washington, DC: Institute for International Economics, 2002).
- “Poor Results and Poorer Policy: A Comparative Analysis of Estimates of Global Inequality and Poverty,” *CESifo Economic Studies* 50 (2004), 85–132.
- Blundell, R. W., M. Browning, and I. A. Crawford, “Nonparametric Engel Curves and Revealed Preference,” *Econometrica* 71 (2003), 205–240.
- Bourguignon, F., and C. Morrisson, “Inequality among World Citizens: 1820–1992,” *American Economic Review* 92 (2002), 727–744.
- Bronars, S. G., “The Power of Nonparametric Tests of Preference Maximization,” *Econometrica* 55 (1987), 693–698.
- Chen, S., and M. Ravallion, “How Did the World’s Poor Fare in the 1990s?” *Review of Income and Wealth* 47 (2001), 283–300.
- “Absolute Poverty Measure for the Developing World, 1981–2004,” *Proceedings of the National Academy of Sciences* 104 (2007), 16757–16762.
- “The Developing World Is Poorer Than We Thought, But No Less Successful in the Fight against Poverty,” *Quarterly Journal of Economics* 125 (2010), 1577–1625.
- Deaton, A., “Counting the World’s Poor: Problems and Possible Solutions,” *World Bank Research Observer* 16 (2001), 125–147.
- “Measuring Poverty in a Growing World (or Measuring Growth in a Poor World),” this REVIEW 87 (2005), 1–19.
- Deaton, A., and A. Heston, “Understanding PPPs and PPP-Based National Accounts,” *American Economic Journal—Macroeconomics* 2:4 (2010), 1–35.
- Deininger, K., and L. Squire, “A New Data Set Measuring Income Inequality,” *World Bank Economic Review* 10 (1996), 565–591.
- Diewert, W., “Afriat and Revealed Preference Theory,” *Review of Economic Studies* 40 (1973), 419–426.
- Dowrick, S., and M. Akmal, “Contradictory Trends in Global Income Inequality: A Tale of Two Biases,” *Review of Income and Wealth* 51 (2005), 201–229.
- Dowrick, S., and J. Quiggin, “True Measures of GDP and Convergence,” *American Economic Review* 87 (1997), 41–64.
- Eltető, Ö., and P. Köves, “On a Problem of Index Number Computation Relating to International Comparisons,” *Statisztikai Szemle* 42 (1964), 507–518 (in Hungarian).
- Geary, R. C., “A Note on the Comparison of Exchange Rates and Purchasing Power between Countries,” *Journal of the Royal Statistical Society (Series A)* 121 (1958), 97–99.
- Gerschenkron, A., “A Dollar Index for Soviet Machinery Output, 1927–28 to 1938” (Santa Monica, CA: Rand Corporation, 1951).
- Heston, A., R. Summers, and B. Aten, “Penn World Table Version 6.3” (Philadelphia: Center for International Comparisons at the University of Pennsylvania, 2009).
- Hill, R. J., “Measuring Substitution Bias in International Comparisons Based on Additive Purchasing Power Parity Methods,” *European Economic Review* 44 (2000), 145–162.
- Khamis, S. H., “A New System of Index Numbers for National and International Purposes,” *Journal of the Royal Statistical Society (Series A)* 135 (1972), 96–121.
- Kravis, I. B., A. Heston, and R. Summers, *World Product and Income: International Comparisons of Real Gross Products* (Baltimore, MD: Johns Hopkins University Press, 1982).
- Maddison, A., *Monitoring the World Economy: 1820–1992* (Paris: Organisation for Economic Co-operation and Development, 1995).
- *The World Economy: Historical Statistics* (Paris: Organisation for Economic Co-operation and Development, 2003).
- Neary, J., “True Multilateral Indices for International Comparisons of Real Income,” *American Economic Review* 94 (2004), 1411–1428.
- Ravallion, M., “How Not to Count the Poor? A Reply to Reddy and Pogge,” mimeograph, World Bank (2003a).
- “Measuring Aggregate Welfare in Developing Countries: How Well Do National Accounts and Surveys Agree?” this REVIEW 85 (2003b), 645–652.
- “Dollar a Day Revisited: The 2005 ICP,” *ICP Bulletin* 5:3 (2008), 10–15.
- “Price Levels and Economic Growth,” World Bank policy research paper 5229 (2010).
- Reddy, S. G., and T. W. Pogge, “How Not to Count the Poor” (2003), <http://www.socialanalysis.org>.
- Sala-i-Martin, X., “The World Distribution of Income (estimated from Individual Country Distributions),” NBER working papers 8933 (2002).
- “The World Distribution of Income: Falling Poverty and . . . Convergence, Period!,” *Quarterly Journal of Economics* 121 (2006), 351–397.
- Samuelson, P. A., “Theoretical Notes on Trade Problems,” this REVIEW 46 (1964), 145–154.
- Szulc, B. J., “Index Numbers for Multilateral Regional Comparisons,” *Przegląd Statystyczny* 3 (1964), 239–254 (in Polish).
- UNU-WIDER, “World Income Inequality Database (WIID) V 2.0b May 2007” (2007), <http://website1.wider.unu.edu/wiid/wiid.htm>.
- Varian, H. R., “The Nonparametric Approach to Demand Analysis,” *Econometrica* 50 (1982), 945–973.
- “Non-Parametric Tests of Consumer Behaviour,” *Review of Economic Studies* 50 (1983), 99–110.
- World Bank, *World Development Report 2000/2001: Attacking Poverty* (New York: Oxford University Press, 2000).
- *Globalization, Growth and Poverty* (New York: Oxford University Press, 2002).
- “World Development Indicators (WDI): 2003” 2003, CD-ROM.
- “Global Purchasing Power Parities and Real Expenditures: 2005 International Comparison Program,” (Washington, DC: International Bank for Reconstruction and Development/World Bank, 2008).
- “WDI Online: World Development Indicators” (2009), www.worldbank.org.