Purchasing power parity exchange rates for the global poor

Angus Deaton, Research Program in Development Studies, Princeton University
Olivier Dupriez, Development Economics Data Group, World Bank

November 2009

We acknowledge invaluable assistance and advice from Bettina Aten, Shaida Badiee, Misha Belkindas, Yonas Biru, Shaohua Chen, Tefera Bekele Degefu, Yuri Dikhanov, Erwin Diewert, Alan Heston, D. S. Prasada Rao, Martin Ravallion, Sam Schulhofer-Wohl, Eric Swanson, Changqing Sun, and Fred Vogel.

1. Introduction

The first of the Millennium Development Goals targets global poverty. The numbers that support this goal are estimated by the World Bank, and come from a worldwide count of people who live below a common international poverty line. This line, loosely referred to as the dollar-a-day line, is calculated as an average over the world's poorest countries of their national poverty lines expressed in international dollars. The counts of those living below the line come from household surveys, the number and coverage of which have steadily increased over the years. National poverty lines are converted to international currency using the purchasing power parity (PPP) exchange rates from the various rounds of the International Comparison Program (ICP). These PPPs, unlike market exchange rates, are constructed as price indexes that compare the level of consumer prices across countries.

In the first dollar-a-day poverty calculations, the World Bank (1990) used price indexes for GDP as a whole, but this practice was later improved by the use of price indexes for household consumption. But even this may be misleading if the price indexes for national aggregate consumption are different from those that are relevant for people who live at or around the global poverty line. Price indexes are weighted averages of prices, and both weights and prices could be wrong. The prices collected by the ICP may be different from the prices faced by those at the poverty line, and the expenditure patterns at the poverty line are almost certainly different from the aggregate expenditure patterns in the National Accounts that provide the weights for the usual consumption PPPs. This paper is concerned with the second of these issues, the recalculation of purchasing power parity exchange rates using the expenditure patterns of those at the

global poverty line. We shall refer to these poverty-weighted purchasing power parities as PPPPs or P4s, as opposed to the aggregate weighted PPPs or P3s. We recognize the possible importance of the first issue but our procedures and calculations use the national prices of goods and services collected by the ICP so that our P4 indexes differ from the P3s published by the ICP only in the methods that we use to turn these prices into national price indexes.

Although our objectives are relatively modest, there are substantial theoretical and technical issues to be faced. First, in order to calculate the appropriate weights in each country, we need to identify those who are close to the local currency equivalent of the global poverty line. But to convert the global line to local currency, we need the P4s, so that the P4s and their weights need to be simultaneously calculated. Second, the global poverty line is itself calculated as an appropriate average of local lines converted to international units using the P4s, so that our calculations need to solve simultaneously for weights, price indexes, and the global poverty line. Third, the current standard procedure uses aggregate data from the *national accounts* to calculate the PPPs and the global poverty line in international dollars, but then takes the global poverty line to *household survey* data to calculate the numbers of poor people in each country. In the calculations in this paper, we use household survey data throughout. We use (a) local currency prices (or more accurately "parities" or commodity specific PPPs) for 102 basic headings of household consumption from the 2005 round of the ICP, (b) nationally representative household surveys from 62 poor countries, and (c) national poverty lines in local currency for 50 countries, and combine (a), (b) and (c) to calculate a set of poverty-weighted purchasing power parity exchange rates for consumption, a global poverty line,

and a set of global poverty counts for each country and the world as a whole. The 62 countries for which we have survey data represent 83 percent of the population of the countries included in the global poverty counts; the 50 poverty lines also cover 79 percent of the population of poor countries. Fourth, when calculating P4s, we cannot follow the usual practice with P3s of taking the US as base because there are no households in the US at a poverty line in the vicinity of a dollar a day, so it is not possible to calculate weights. Our calculations use only information from the much poorer countries included in the global poverty count. This has the advantage that prices and expenditure patterns in rich countries have no effect on P4s or on the global poverty count, and that we are not using a "global" poverty line at which much of the (rich) world could not survive.

The paper is laid out as follows. In Section 2 we review the theory of the P4 indexes and the differences between P3s and P4s. We work with three different types of multilateral indexes, the Fisher and Törnqvist versions of the EKS index, and the weighted country product dummy index. We explore two different methods for solving the simultaneity. We show that the P3 and P4 indexes for any pair of countries will differ according to the cross-commodity correlation between relative prices and income elasticities. If food is relatively expensive in poor countries, this will raise the P4 relative to the P3 for a poor country relative to a rich country, but these differences will be moderated between poor countries as a group. We also discuss the construction of standard errors for our price indexes. One concern is with the sample size of some of our household surveys, so that we need to ensure that using samples, as opposed to populations, does not affect the precision of the estimates. Another concern is related to the fact that, in a world where relative prices are different in different countries, different

index number formulas give different answers, and we develop a standard error concept that captures the degree of uncertainty from this cause.

Section 3 discusses practical issues. We discuss how the ICP constructs the prices for the basic heads of consumption, and how we need to modify those procedures. We discuss the matching of consumption categories in the household surveys with the basic headings of consumption in the ICP and note that there are several categories—rent and health being perhaps the most important—that are not adequately represented in the surveys. Beyond that, some surveys contain imputations for the use value of durables, as opposed to expenditures on those items in the national accounts and the ICP. As a result, even when we calculate P3s as opposed to P4s, our estimates will not coincide with those in the ICP. A final practical issue is that, for some countries, the ICP collected only urban prices, and we have good evidence from many countries that urban prices are higher than rural prices, so that an adjustment is necessary.

Section 4 presents our results. We present our estimates of P3s and P4s for 62 of the countries included in the global poverty counts and compare them with the P3s from the ICP itself. Perhaps our major conclusion is that, provided we use household survey data in both calculations, the reweighting to a poverty basis makes little difference, so that our P3s are close to our P4s. However, our P3s are somewhat further away from the P3s in the ICP, in part because of our different aggregation procedures (definitions of the indexes), and in part because survey–based estimates of aggregate expenditure patterns often differ from the those presented in the national accounts. As is often the case in poverty work, data discrepancies are more important than definitional or conceptual issues. We also use our P4s to calculate poverty counts, by region and for the world as a

whole; although our P4s are close to the official P3s, our poverty count is a good deal lower than the official count because of the way that we construct out global poverty line.

2. Poverty-weighted purchasing power parity exchange rates: theory

Purchasing power parity exchange rates are multilateral price indexes designed to summarize price levels in each of a group of countries. In this paper, we are interested in price indexes for household consumption, and wish to depart from the standard practice of calculating indexes for aggregate national consumption. Instead, our aim is to calculate indexes using weights for people that are at, or at least close to, the global poverty line.

We start with notation. We have M countries, labeled using the index c. In each country, there is a vector of prices for N items of consumption, labeled using the index n, so that p_n^c is the price of good n in country c. Associated with those prices is a pattern of consumption, which we shall typically measure in terms of the shares of the budget devoted to each good, denoted s_n^c. The sum of these non-negative budget shares over n is unity for each country c, so that they can be thought of as weights. They are defined as the expenditure on each good divided by the total expenditure on all goods and services. Each household has a set of budget shares, and the economy as a whole has budget shares defined as aggregate expenditure on each good divided by aggregate total expenditure on all goods. We shall distinguish these as necessary.

Throughout the work described in this paper, we shall assume that the prices are the same for all consumers in the country, and we will use price data on 102 "basic headings" of consumption collected by the ICP. The expenditures and prices of these basic headings are themselves aggregates of the thousands of narrowly defined goods and services

whose prices are collected in the ICP; in our work, we do not go below the parities of the basic heading aggregates, which we treat as our underlying prices. The difference between what we do and the standard practice is in the treatment of the budget shares or weights. In the national accounting treatment of the ICP, the weights are the shares of aggregate national expenditure spent on each good whereas, in our treatment, the weights are calculated from household surveys, and are defined as an average of budget shares for households close to the global poverty line. The global poverty line is itself defined as an average of local poverty lines expressed in PPP terms, and we shall show how to measure the price indexes and the global poverty line simultaneously. For the moment, we assume that we know the line, and that we have calculated the poverty-line budget shares for each country.

There are two different types of PPP indexes that we shall compute, the Elteto-Köves-Sculc (EKS) type, and the weighted country-product-dummy (CPD) type. EKS indexes begin from a set of superlative indexes (Diewert, 1976) calculated for each pair of countries. We work with two familiar superlative indexes, the first of which is the Törnqvist index, defined as

$$\ln P_T^{cd} = \frac{1}{2}\sum_{n=1}^{N}(s_n^c + s_n^d)\ln\frac{p_n^c}{p_n^d} \tag{1}$$

Note that we adopt the convention that the base country, here country *c*, comes first in the superscript on the index, followed by the comparison country, here *d*. The Törnqvist index is thus a weighted geometric average of the price relatives of each good, with the weights the average of the two budget shares in *c* and *d*. We leave the precise definition of the budget shares for later, but (1) will apply whatever budget shares we use.

The second familiar index is the Fisher ideal index, defined as the geometric mean of the Paasche index and the Laspeyres index so that, in logarithms,

$$\ln P_F^{cd} = 0.5 * \ln\left[\sum_{n=1}^{N} s_n^c \frac{p_n^d}{p_n^c}\right] - 0.5 * \ln\left[\sum_{n=1}^{N} s_n^d \frac{p_n^c}{p_n^d}\right] \tag{2}$$

The first term in brackets on the right-hand side is the Laspeyres index for *d* relative to *c*, while the second term in brackets is the Laspeyres for *c* relative to *d*, which is identical to the reciprocal of the Paasche for *d* relative to *c*. The log Fisher and Törnqvist indexes in (1) and (2) give us an *M* by *M* matrix of index numbers comparing every country with every other country. In practice, a matrix of price indexes is less useful than a vector of price levels, one for each country relative to a numeraire country, with each representing a purchasing power version of exchange rates. In order to compress the information in (2) into this form, the matrix is converted into a set of international PPP exchange rates by applying an adjustment first proposed by Gini (1924), and later rediscovered, so that it is now referred to as the EKS procedure.

If we write *B*, typical element b^{cd}, for the *M* by *M* matrix of unadjusted log price indexes (2), the EKS log PPPs are given by

$$a^c = \frac{1}{M}\sum_{j=1}^{M}(b^{1j} + b^{jc}) \tag{3}$$

where country 1 is the arbitrarily designated numeraire country. In terms of the original prices the EKS PPP price index for *c* in country 1's units is

$$P_F^c = \left(\prod_{j=1}^{M} P_F^{1j} P_F^{jc}\right)^{\frac{1}{M}} \tag{4}$$

for the EKS-Fisher, with an identical formula, with *T* replacing *F*, for the EKS-Törnqvist. Each index inside the brackets is the price level of *c* relative to 1 computed via country *j*,

so that the EKS index comes from taking a geometric average of these indexes over all possible intermediate countries.

We shall also work with PPP indexes constructed according to the weighted country product dummy method; an unweighted version of this traces back to Summers (1973), with the weighted version developed by Prasada Rao, see for example Selvanathan and Rao (1994), and Rao (1990, 2005) . If the law of one price were true and there were perfect price arbitrage in goods and service across countries, prices would differ only in currency units, so that we would be able to write

$$\ln p_n^c = \alpha^c + \beta_n \tag{5}$$

where α^c is the logarithm of the value of country *c*'s currency relative to country 1, for which $\alpha^1 = 0,$ and β_n is the price of good *n* in country 1, which is also the price of good *n* in all countries, up to unit scaling. If (5) were true, PPP exchange rates would be equal to market exchange rates. In reality, we can construct a set of price indexes that approximate the structure (5) by *projecting* actual prices on to a set of country and product dummies by running a weighted regression of the form

$$\ln p_n^c = \alpha^c + \beta_n + \varepsilon_n^c \tag{6}$$

in which the weights are the budget shares of each good in each country, s_n^c. The intuitive argument for the budget-shares weights is the same as for other price index calculations, that goods with large (small) budget shares should count more (less) in the calculations.

For future reference, the formula for the weighted CPD is

$$\hat{b} = (X'SX)^{-1} X'Sy \tag{7}$$

where X is an MN by $N+M-1$ matrix of ones and zeroes, with $N-1$ columns for the commodities, $M-1$ columns for the countries, plus a constant and rows corresponding to the vector y, which is the "stacked" vector of log prices, N for each of M countries. The S matrix is a diagonal matrix with the budget shares on the diagonal, N for each country. The element of the estimated parameter b corresponding to the country dummies are the estimates of the log of the weighted CPD-PPP exchange rates for each country in terms of country 1. Note that, although (7) can be thought of as a generalized least-squares estimator, the weighting matrix S is included for substantive reasons to do with the importance of each good in spending patterns, and not because of any supposed relationship between the budget shares and the variances of the error terms in (6). Indeed, (6) should not be thought of as a "true" model of the data generating process; rather (6) and (7) should simply be regarded as a convenient device for projecting the log prices on to country and commodity effects in a metric that recognizes the different importance of each commodity in the budget.

We make no use of the Geary-Khamis (GK) system of PPPs as used, for example, in the Penn World Table but which has disadvantages for poverty work. The most important of these is that the GK method prices all goods at world prices that are quantity weighted averages of individual country prices, so that countries with the largest physical volume of consumption of a good gets greatest weight in the construction of the composite world prices. The use of such prices has the effect of overstating the level of consumption—and underestimating poverty—in the poorest countries, the Gershenkron effect.

The EKS and CPD formulas allow us to calculate a set of PPPs given budget shares and prices for each country. In our calculations for poverty-weighted PPPs, we use the

budget shares for households at or near the global poverty line. This distinguishes our PPPs (P4s) from the consumption PPPs (P3s) from the ICP, in which the budget shares are the shares of aggregate consumers' expenditure on each good in the aggregate of consumers' expenditure in total. If s_n^{ch} is the budget share on good *n* by household *h* in country *c*, the aggregate budget shares that go into the ICP indexes can be written

$$\tilde{s}_n^c = \frac{\sum_{h=1}^{H} x^{ch} s_n^{ch}}{\sum_{h=1}^{H} x^{ch}} \tag{8}$$

where x^{ch} is the total expenditure of household *h*. Indexes using weights such as (8) are referred to as plutocratic indexes, Prais (1959), because the budget share of each household is weighted by total expenditure and those who spend more are counted more.

The weights that we shall use for the poverty PPPs are, not (8), but

$$\overline{s}_n^c(z^c) = E[s_n^{ch} \mid (x^{ch} / n^{ch}) = z^c] \tag{9}$$

where n^{ch} is household size and z^c is the poverty line in local currency, so that according to (9), the budget shares for poverty weighting are the average budget shares of households at the poverty line. This is indexed by the country *c* because it is the value in local currency of the global poverty line. Note that the averages in both (8) and (9) include the budget shares of all households, even if those who do not purchase a good whose budget shares are zero. A household who buys nothing of good *n* is unaffected by changes in its price, and this weight needs to be counted in the overall index. This is also relevant because in many surveys, especially around the poverty line, some goods are bought by only a few households, and the precision of the estimates will depend on the

total number of households (or the total number near the poverty line), not on the number who purchase.

The global poverty line is expressed in international currency—most famously the dollar a day line—while the calculation of the budget shares from the surveys in (9) requires that the line be expressed in local currency at its purchasing power equivalent. In consequence, the expenditure weights used to calculate the price indexes require that we know the price indexes before we start. We propose two methods for dealing with this issue. The first allows the calculation of an exact, one-step, solution that works only for the Törnqvist index and that requires that the Engel curves in each country have a specific functional form. The second is an iterative procedure that uses the first method to provide starting values.

For the exact method, we start from the two-country case. Suppose that the global poverty line in country 1's currency is *z*. The budget shares in each country are a function—among other things—of household total per capita expenditure (PCE) $\tilde{x} = x / n$, which we write as $s_n^c(\tilde{x})$ for good *n* in country *c*, with the function interpreted as the expected budget share for households with PCE of *x*. The equation we need to solve for the relevant Törnqvist PPP is

$$\ln P_T^{12} = \frac{1}{2}\sum_{n=1}^{N}\left[s_n^1(z) + s_n^2(P_T^{12} z)\right] \ln \frac{p_n^2}{p_n^1} \tag{10}$$

so that the budget shares for the index are at the global poverty line in both countries. Suppose that the budget shares in each country are linear functions of the logarithm of total expenditures, a functional form that often fits the data well, and that is consistent with choice theory, see for example, Deaton and Muellbauer (1980, Chapter 3.)

$$s_{nh}^{c} = \xi_{0n}^{c} + \xi_{1n}^{c} \ln \tilde{x}_{h} + v_{nh}^{c} \tag{11}$$

where c is the country, here 1 or 2, v_{nh} is a disturbance term, and ξ_{0n}^{c} and ξ_{1n}^{c} are commodity- and country-specific parameters. For each country, the ξ_{1n}^{c} parameters add to zero over all the goods in the budget, while the ξ_{0n}^{c} parameters to one. If we substitute the conditional expectation of (11) into (10), the poverty-line Törnqvist index can be written

$$\ln P_{T}^{12} = \frac{\sum_{n=1}^{N}\left(\xi_{0n}^{1} + \xi_{0n}^{2} + (\xi_{1n}^{1} + \xi_{1n}^{2}) \ln z\right) \ln \frac{p_{n}^{2}}{p_{n}^{1}}}{2 - \sum_{n=1}^{N} \xi_{1n}^{2} \ln \frac{p_{n}^{2}}{p_{n}^{1}}} \tag{12}$$

which is in closed-form and can be calculated directly from the prices, the budget shares, and the global poverty line.

The M–country extension of (12) is straightforward. Assuming the same set of Engel curves (11), the logarithm of the Törnqvist index for j in terms of i is written

$$b^{ij} = \frac{1}{2} \sum_{n=1}^{N} \left((\xi_{on}^{i} + \xi_{on}^{j}) + (\xi_{1n}^{i} + \xi_{1n}^{j}) \ln z + \xi_{1n}^{i} a^{i} + \xi_{1n}^{j} a^{j} \right) \ln \frac{p_{n}^{j}}{p_{n}^{i}} \tag{13}$$

where a^{i}, from (3), is the Törnqvist-EKS PPP-exchange rate for country i in terms of country 1. This can be rewritten in the form

$$b^{ij} = \psi^{ij} + \theta^{ij} a^{i} - \theta^{ji} a^{j} \tag{14}$$

where the definitions of the new terms can be read off from (13). Given the relative prices, the coefficients of the Engel curves, and the global poverty line, the quantities ψ^{ij} and θ^{ij} are known. Equation (3) also links the EKS–Törnqvist PPPs to the pairwise Törnqvist indexes b^{ij} so that, if we combine (3) and (14), we reach

$$a^{i}\left(1 + \frac{1}{M} \sum_{j=1}^{M} \theta^{ij}\right) + \frac{1}{M} \sum_{j=2}^{M} (\theta^{j1} - \theta^{ji}) a^{j} = \frac{1}{M} \sum_{j=1}^{M} (\psi^{1j} + \psi^{ji}) \tag{15}$$

where we have used the fact that $a^1 = 0.$ Equation (15) is a system of $M-1$ linear equations in the $M-1$ unknown EKS–Törnqvist P4-indexes under the assumption that the Engel curves take the form (11).

In general, none of the EKS–Törnqvist, EKS–Fisher, or weighted CPD P4 index has a closed-form solution. Instead, we start from the global poverty line converted to local currencies using the Törnqvist approximation (or some other set of PPPs, such as the consumption PPPs from the ICP), calculate a set of budget shares for households at or near those poverty lines in each country, which are used to calculate a new set of poverty-weighted PPPs. At the next iteration, these are applied to the global poverty line instead of the original starting values, and so on.

We calculate "near the line" budget shares by computing a weighted average of the budget shares in the sample with weights that are largest at the poverty line, and decline as we move away from it. Define the weight $\omega_{h\tau}^c(z)$ for household *h* in country *c*

$$\omega_{h\tau}^c(z) = \frac{1}{\tau} K\left(\frac{\ln \tilde{x}_h - \ln z - \ln a^c}{\tau}\right) \tag{16}$$

where z is the global poverty line in numeraire currency. The function *K*(.) is a kernel function that integrates to unity, is non-negative, symmetric around zero, and decreasing in the absolute value of its argument and τ is the bandwidth, declining in the sample size, that is ideally set to optimize the trade-off between bias (too large a bandwidth with many households far from the line) and variance (too small a bandwidth but all households near the poverty line).

In general, it is not possible to guarantee that there exists a unique solution for the set of poverty-weighted PPP indexes. However, we know that uniqueness is guaranteed for

the EKS–Törnqvist when the Engel curves satisfy (11). It is also straightforward to show that in the case where all countries have the same tastes, and the price indexes are cost-of-living indexes, there is a unique solution. Given that both the Fisher and Törnqvist indexes are superlative indexes, this result would be useful if we could accept the position that there is no international heterogeneity of tastes. Further investigation of the issue is contained in Deaton and Schulhofer-Wohl (2009).

In order to interpret our results, it is useful to investigate the differences between the various indexes, between different types, EKS–Fisher, EKS–Törnqvist, and weighted CPD, and between indexes that use poverty weights versus those that use aggregate weights. The two country Törnqvist approximation (12) can be used to make the main points. If the budget shares do not vary with total household expenditure, the parameters ξ_{1n}^c in (12) are zero, so that the term involving z in the numerator of (12) and the second term in the denominator are both zero. In this case (12) is simply the P3 Törnqvist index, because the ξ_{0n}^c parameters are the averages of the budget shares, and because the budget shares do not vary with income, they are also equal to the aggregate weights so that (8) and (9) coincide.

More generally, the difference between the poverty-weighted and plutocratic Törnqvist indexes can be written

$$\ln P_T^{12} - \ln \tilde{P}_T^{12} = 0.5\sum_{n=1}^{N}[\xi_{1n}^1(\ln z^1 - \ln y^1) + \xi_{1n}^2(\ln z^2 - \ln y^2)]\ln\frac{p_n^2}{p_n^1} \qquad (17)$$

where z^1 and z^2 are the two local currency poverty lines, and y^c is an (entropy) inequality adjusted measure of mean expenditure

$$\ln y^c = \sum_h \left[\frac{x_h^c \ln x_h^c}{\sum_h x_h^c} \right] \tag{18}$$

and where y^c is measured in local prices. These equations tell us that, if the effects of income on the budget shares, as measured by the ξ_{1n}^c parameters, are orthogonal, for each country, to the logarithms of the price relatives, the plutocratic and poverty-weighted indexes will be the same. When these orthogonality conditions fail, the plutocratic and poverty-weighted indexes will differ by an amount that depends on the correlation between the ξ_{1n}^c's and the relative prices, on the inequality-adjusted levels of living in the two countries, and on the poverty line.

To illustrate with an important case, if we are comparing a rich(er) country with a poor(er) country, and if food in both is mostly traded, then food will be relatively expensive in the poor country, as is typically the case. Suppose that there are only two goods, food *f*, and non-food *n*, and that the Engel curve parameters ξ_{1n} are the same in both countries. The food parameter is typically estimated to be around –0.15, so that the non-food parameter is 0.15. Then the numerator of (17) simplifies to

$$\xi_{1f} \ln \sqrt{\frac{z^1 z^2}{y^1 y^2}} \ln \left(\frac{p_f^2}{p_n^2} \Big/ \frac{p_f^1}{p_n^1} \right) \tag{19}$$

which is positive if food is relatively more expensive in the poor country, and if the poverty lines are less than inequality-adjusted mean expenditure in both countries. In this example, the P4 index for the poor country relative to the rich country will be *higher* than the corresponding P3 index, essentially because the food share is declining in income and the relatively higher food price gets more weight in the P4-index than in the P3-index.

The size of the effect will be larger the larger the Engel effect, and the larger the distance between the poverty lines and inequality-adjusted mean expenditures in both countries. It is a good deal harder to think of any such systematic effects between countries at similar levels of development which, as we shall see, is the relevant case here where we calculate P3s and P4s for a set of relatively poor countries.

The above argument is specific to the Törnqvist and to the two country case. But the argument about the correlation between Engel patterns and the structure of relative prices is clearly a general one, and should serve as a rough guide to the way in which we would expect P4 indexes to differ from P3 indexes. The extension to multiple countries is harder to derive formally, but practical experience has been that the EKS adjustment of the matrices of Fisher and Törnqvist indexes is typically not very large, so that the final index is likely to be dominated by the pairwise indexes, not by the final EKS adjustment.

We shall calculate three different indexes, and it is useful to understand something about how they might be expected to relate to one another. The EKS indexes come from an adjustment of a set of country pairwise superlative indexes. In consequence, if as usually seems to be the case in practice, the EKS adjustment (3) makes relatively little difference, so that $b^{cd} \approx a^c - a^d$, the Fisher and Törnqvist EKS indexes will share the properties of their parent superlative indexes. There is no similar argument for the weighted CPD index. Another useful point comes from thinking of equation (6), which we used to define the weighted CPD index, as an approximation. According to this, we would think of the international structure of prices as being approximated by a common set of relative prices, scaled up by a set of purchasing-power converters, one for each country. If we substitute (6) into the formulas for the various different indexes, it is easy

to show that, ignoring powers of ε_n^c beyond the first, the Fisher and Törnqvist indexes are identical, as will be the EKS indexes derived from them. The approximation for the weighted CPD index is different, so we might expect the two EKS indexes to be closer to one another in practice. In the empirical results in section 4, we shall consistently find that this to be true.

So far, we have assumed that we know the global poverty line in the numeraire country as, for example, when the line is a dollar a day per capita in international PPP dollars. However, the dollar a day line is itself revised with revisions in P3s, see Ravallion, Chen, and Sangraula (RCS) (2009) for the latest treatment, so we need a procedure to adapt the international line to our calculated P4s. We cannot use the dollar standard, because the US is not one of our countries, so we (arbitrarily) choose India, and set our global line in "world rupees"; the choice of India makes no substantive difference, and the results would be the same up to scale for any other country. We consider three variants. The first variant, and our baseline case, calculates a global poverty line from 50 countries that are included both in our set of household surveys and in the compilation of local poverty lines in RCS. At each iteration in the P4 calculations, we convert these 50 lines to world rupees, and take a weighted average using as weights the numbers of people below the line in each of the countries. The second variant is the same as the first, but with the 50 local poverty lines multiplied by two before we start; this is similar in spirit to looking at one and two dollars a day. Our third variant follows RCS and calculates the international line as the simple average of the world rupee value of the local poverty lines of Chad, Ethiopia, Gambia, Ghana, Malawi, Mali, Mozambique, Nepal, Niger, Rwanda, Sierra Leone, Tajikistan, Tanzania, and Uganda. (RCS also

include Guinea-Bissau, for which we lack survey data.) Deaton (2010) discusses the advantages and disadvantages of each of these different procedures.

A final issue for this section is the calculation of standard errors for our P3 and P4 estimates. Our calculations use not national accounts numbers, but household surveys whose sample sizes vary from country to country. Sample sizes are further restricted when we focus on households close to the poverty line. The calculation of these sampling standard errors is straightforward in principle; all of the P4s presented above are functions of sample means from the surveys, whose designs—sample sizes, weighting, stratification—we know. The formulas are derived in detail in Deaton and Dupriez (2009), and can be implemented using any software that handles complex survey design.

We also provide a second kind of standard errors which we refer to as the "failure of arbitrage" standard errors. These come from the following conceptual experiment. Suppose that we write the price of good *n* in country *c* in the form (6) in which the logarithm of price is the sum of a country effect, a commodity effect, and an error. In a world of perfect arbitrage, where relative prices were the same in all countries, and absolute prices differed only according to the currency unit, the error terms in (6) would be zero, and the α^c would be the logarithms of the PPPs, of the exchange rates, or of any reasonable index of prices in the country. Because perfect arbitrage does not hold, the ε_n^c are not zero, and different index number formulae will give different answers. It is this variability across indexes that is captured by the "failure of arbitrage" standard errors. This measure of model uncertainty is similar in concept to the use of the "Paasche-Laspeyres spread," another measure of the extent to which different price formulas give different answers when relative prices differ across countries.

In calculating our "failure of arbitrage" standard errors, the conceptual experiment is one in which we think of ε_n^c as drawn repeatedly, which generates stochastic prices according to (6), which are then combined with non-stochastic expenditure weights to generate stochastic P3s and P4s whose standard errors are calculated. Note that these standard errors are conditional on the budget shares which we take as fixed. It is easy to imagine an alternative set of standard errors which models the dependence of the weights on the prices, for example through a cross-country model of consumer behavior. We do not consider that extension here, in large part because we do not want to commit to any such model, instead regarding the failure of arbitrage standard errors as descriptive measures of the dispersion of the ε_n^c, not directly, but through the PPP indexes.

Once again, the formulas are developed in Deaton and Dupriez (2009). To illustrate briefly, consider the weighted-CPD P3 and P4 indexes. The CPD indexes are estimated by running the generalized least squares regression (8), and an estimate of the variance covariance matrix of the estimated parameters can be obtained from

$$\widehat{V}(\widehat{b}) = (X'SX)^{-1}(X'S\Sigma SX)(X'SX)^{-1} \tag{20}$$

where X is the matrix of country and product dummies, S is a diagonal matrix of the budget share weights, and Σ is an estimate of the variance-covariance matrix of the ε_n^c, the deviation of the log prices from perfect arbitrage. In practice, we estimate Σ by a diagonal matrix containing the squares of the estimated residuals from the CPD model. The derivations for the EKS indexes are more complex, but follow the same principles.

3. Practical issues: linking ICP prices to household survey data

We are now in a position to discuss how to bring together the prices of goods and services from the ICP and the budget weights from the household surveys. There are some immediate differences between the two projects. First, the ICP covers all of the countries in the world, at least in principle, while our interest is confined to the countries that are included in the global poverty count. As we shall see, this necessitates some prior screening and processing of the ICP price data. Second, not all of the relevant countries in the ICP have household surveys, and some do not allow them to be used for poverty-related analysis. Third, the surveys that we have were not collected for the purpose of calculating international price indexes. In particular, the categories of consumption for which we have data are not uniform across countries, and none match exactly the list of consumption goods that is used for the ICP itself, some of which are not covered in the surveys at all. We discuss each of these issues in turn.

At its heart, the ICP is a large-scale price collection effort in which a list of commodities is priced in many countries. In practice, it is impossible to use a single list for all countries of the world, and for this and for management reasons, the 146 countries that were included in the 2005 round were broken up into six geographic regions. At a first stage, each region carried out its own regional calculations in which PPP indexes were calculated for all of the countries in each region, with a separate numeraire currency in each region. At a second stage, these regional estimates were linked to give a global set of PPPs with the (international) US dollar as the unit of account. At the first stage in each region, the prices for the detailed regional list in each country are combined to give prices for 155 “basic headings” of GDP, 110 of which are items of “individual consumption expenditures by households.” These are then linked through a set of “ring” countries,

strategically placed in each region, to give a global list of basic heading parities in a single numeraire currency; the process was developed by Diewert (2008), which contains a full account, see also Hill (2007a, 2007b). Deaton and Heston (2008) explain the procedure in more detail and discuss some of its strengths and weaknesses.

For the calculations here, we recalculate the global list of parities for basic heads but with the OECD region dropped because we want our calculations to exclude price data from the rich countries. Our global P4s are developed entirely from information from the countries whose poverty is being measured, and neither the total number of global poor, nor of the globally poor in any poor country, should depend on commodity prices or expenditure patterns in rich countries. In practice, this change makes very little difference, and the prices we use for each basic heading in each country are almost identical to those used by the ICP. Given those prices, and the 62 ICP non-OECD countries for which we have survey data, we calculate our P3 and P4 indexes treating all countries simultaneously irrespective of their regional affiliation.

When the survey categories are finer than the basic headings for consumption in the ICP, they can be aggregated up to match. The harder case is when the categories are larger in the survey than in the ICP, or are neither larger nor smaller, but different. For example, one basic head in the ICP consumption is "butter and margarine;" a survey might have these two separate, or part of a larger group "butter, margarine, and edible oils," or have two categories, one of which contains butter together with other items, and one of which contains margarine together with other items. In the two last cases, our procedure is to aggregate the survey categories until we have a category that contains multiple whole basic headings, and then to split the aggregate according to the

proportions in the national accounts on a household by household basis. Following the same example, if we have a survey category "butter, margarine, and edible oils" and if the country's national accounts show that, in aggregate, two-thirds of the category is edible oils, we then go through the survey data, household by household, and allocate two-thirds of each household's recorded expenditure to edible oils, and one third to butter and margarine. There are clearly other and potentially more sophisticated ways of synchronizing the two lists, some of which might be worth experimental calculations. However, the example of butter and margarine was chosen to illustrate a typical case. All of the surveys used here have many categories of consumption, and there is no case in which we were forced to allocate large groupings, such as cereals, let alone all food.

In all cases, we used the latest national household survey that was available to us. In the worst cases (Argentina and Djibouti from1996 and Burundi from 1998), weights calculated from the survey were almost a decade older than the ICP prices (2005). All of the other surveys used here are post 2000, with 2003 the modal year; the countries, survey names, and year of data collection are listed in Appendix A4 of Deaton and Dupriez (2009). While it would be ideal to be able to match expenditure weights to the year of survey prices, we would expect the expenditure patterns—especially those of the poor— to change slowly enough that even a lag as long as a decade is unlikely to invalidate the procedure. Indeed, most statistical offices around the world construct their domestic consumer price indexes with weights that are several years (in extreme cases several decades) older than the prices themselves.

There are a number of cases where consumption items that are basic headings in the ICP do not appear in the survey. Indeed, there is considerable diversity in survey

questionnaires and methodology. The number of consumption items covered in questionnaires varies from 39 in Djibouti (recall method, with 64 out of the 105 basic headings omitted) to 6,927 in Brazil (diary method, with only 7 basic headings not covered). On average, 23 of the 105 basic headings are "missing" in survey questionnaires. In most cases, these are basic headings that represent very limited consumption shares (e.g., animal drawn vehicles). It is clear that there is an urgent need to improve and harmonize practices of household consumption measurement.

It is useful to separate items that are indeed consumed, but are not collected in the survey, from items that are not consumed but still appear in the ICP lists. The most important example of the former is owner-occupier rents. Such imputed flows are rarely collected directly (though in places where there is an active rental market, it is sometimes possible to ask owners how much their home could be rented for), but are imputed ex post from housing characteristics weighted up according to the coefficients in a hedonic regression estimated on the (selected) subset of rented houses. This method is probably good enough to give an average for the national income accounts, but we doubt that it gives adequate answers at the individual level, and we were not successful in calculating satisfactory estimates to add back into our surveys. One major concern with any attempt to do so is that rental markets are mostly urban, so that a hedonic regression will primarily reflect the value of housing amenities in towns and cities. To take those coefficients and use them to impute rents to rural housing runs the risk of attributing consumption to the poor that bears little relationship to the real rental value of their homes. The situation is further compromised by the fact that, in many of our surveys, we do not have adequate documentation of how the rental category was constructed. Given

this, and some unsatisfactory early experiments, we eventually dropped the rental category from all the surveys, so that our P3s and P4s exclude this category; note that "dropping" a category is equivalent to assuming that its P3 or P4 is the same as the overall P3 or P4 for the country. This is clearly unsatisfactory but is probably the best that can be done, especially once we recognize that the ICP parities for this category are also problematic, see Deaton and Heston (2008) and Deaton (2010) for discussion.

An even more extreme case is financial intermediation indirectly measured (FISIM). According to current national accounting practice, the profits of banks and insurance companies which, in competitive markets, would be equal to the value of financial intermediation and risk-bearing services to their customers are added into the estimates of consumption by households. Once again, these items do not show up in the surveys. While we can imagine imputing FISIM to survey households according to some formula, we have chosen not to do so, in part reflecting our skepticisms about the extent to which households around the global poverty line receive much benefit from these services.

There are also a number of items that are (almost) never represented in the surveys, and which in some cases never appear in the ICP price surveys, including purchases of narcotics and prostitution, as well as "purchases by non-residential households in the economic territory of the country." Together with rent and FISIM, we drop these items from the lists. A number of other expenditure items are also excluded, namely purchases of animal drawn vehicles, the maintenance and repair of major durables used for recreation and culture, and purchases by residential households in the rest of the world (though some of these items are probably included in other basic headings.) After all of

these exclusions, our calculations are based on 102 out of the 110 consumption basic headings in the ICP.

There are also items that are included in the ICP but are not purchased in some countries. Two notable examples are pork and alcohol in Muslim countries. These cases are different from FISIM, prostitution, or narcotics, in that there are also no prices for these items in the countries where they are not consumed. We do not want to drop these items, however, because there are valid observations on both prices and expenditures for the majority of the countries in the groups, and we do not want to discard that information. For such cases, our procedure is to impute the missing price using the CPD-regressions (6) so that, for example, we impute a price for pork in Bangladesh using the country-effect for Bangladesh (which essentially gives us the exchange rate for Bangladesh) and the "pork effects" from the other countries, which give us a typical relative price for pork. We then leave the item in the survey expenditure files, but assign zero expenditure to all households.

One aspect of the surveys that cannot be defended is measurement error. There are good studies for a number of countries that compare national accounts and survey estimates of comparably-defined items, and that frequently find large differences. For example, Triplett (1997) has found such differences for the United States, even for items that are almost certainly well-measured in the national accounts. Studies in India tend to favor the accuracy of the survey estimates over those from the national accounts, at least for food and apart from some special cases, Kulshesthra and Kar (2005). Note that we are not concerned here with the increasing divergence in many countries between total expenditures in the surveys and the national accounts, documented for example in Deaton

(2005). That discrepancy is important for the measurement of poverty (and of GDP), but price indexes are invariant to the scale of consumption and depend only on its distribution. Unfortunately, the plausible accounts of the survey error—selective non-response by the richest or poorest households, or item-based non-response—will also affect the distribution over commodities. In consequence, differences in indexes—even aggregate plutocratic indexes—according to whether they are constructed with national accounts or survey weights will reflect both deliberate choices about the definition of goods, and accidental choices that come from poorly understood measurement errors.

Another important issue is the treatment of China. China collects household survey data from both rural and urban households and publishes summary tables annually in the Statistical Abstract of China. However, the household level data were not made available to us for this work. Adding China to the list of countries without data is unattractive given its importance in the poverty calculations, and to avoid this we use the published data in a way that allows us to estimate the pattern of expenditures for Chinese households at various levels of household per capita expenditure, essentially by interpolating using the information in the published tables. An account of our procedures is given in Appendix A2 of Deaton and Dupriez (2009).

A final issue in matching ICP prices to the surveys is the treatment of rural and urban sectors. All of our surveys are nationally representative and cover both rural and urban households. In contrast, the ICP collected only urban prices in a number of countries, including most of Latin America, but also in China while, in India, urban outlets were overrepresented in the price surveys. For the urban only countries, we need a measure of the price of consumption in rural relative to urban, and for this we follow Chen and

Ravallion (2008) and use the ratio of rural to urban poverty lines in those countries. While it is a big assumption that the ratio of the poverty lines correctly measures the relative price levels, there is no other obvious source of such information, and some correction is necessary. For countries where the adjustment is made, we adjust our surveys prior to the calculations by converting all household expenditures to urban prices by scaling up per capita household expenditure for each rural household by the ratio of the urban to rural poverty line. Once this adjustment is made, the sectors are ignored, and the survey treated as a single national sample to which the global poverty line, converted at the urban PPP, can be applied to calculate expenditure weights and counts of the numbers in poverty. India is treated somewhat differently first, to take account of the fact that, although the ICP collected both urban and rural prices, the former were over-represented, and second, to recognize that the ratio of official urban to rural poverty lines is implausibly high, and has long been suspected to be the result of a computational error, Deaton (2003). Deaton and Dupriez (2009, Appendix A1) details the Indian calculations.

4. Results

4.1 P3 price indexes from surveys and national accounts

Table 1 shows our calculations of the aggregate (or plutocratic) purchasing power parity exchange rates for household consumption together with those from the ICP. There are 62 countries, and they are listed regionally, Asia first, then South America, Western Asia, and Africa. The ICP numbers in the first column come from the ICP final report, World Bank (2008a), and relate to "individual consumption expenditures by households." Our own calculations in this table, with two calculations each for EKS-Fisher, EKS-

Törnqvist, and weighted CPD, use both surveys and national accounts, so that both sets of weights relate to aggregate national purchases, with one estimated from the surveys and one estimated directly from the national accounts. If the survey and national accounts consumption data were consistent, and had the same coverage of goods and services, the two calculations would give the same results. The ICP estimates in the first column are a subset of the global estimates that come from the global parities for each basic heading, which were constructed differently from our numbers, see the discussion in Section 3 above. Our calculations, for both national accounts and survey-based aggregate weights, treat all 62 countries symmetrically in a single calculation. We are also using parities for the basic heads that were recalculated without data from the rich countries, see Section 3 above, though this made almost no difference in practice.

In Table 1 all of the P3 exchange rates are divided by the market exchange rates listed in World Bank (2008a) so that these numbers can be interpreted as the "price of consumption" in each country. This measure allows us to express all of the indexes in the same units, unobscured by differences in the "size" of currencies which leads to PPP rates that can range from 1000 to 0.001, and eases formal comparison between the indexes. The base country is India, so that all Indian figures are unity. For other countries, if the price of consumption is less than one, the P3 exchange in terms of rupees is lower than the market exchange rate in rupees, so that a rupee converted at the market exchange rate will buy more consumption than it will in India. According to the ICP numbers in column 1, Fiji (2.59), Cape Verde (2.49), Gabon (2.38), and the Maldives (2.15) have the highest consumption price levels among these countries—for comparison, the figure for the US is 2.83—and only Tajikistan (0.84), Kyrgyzstan (0.89), Bolivia

(0.90), Ethiopia (0.90), Paraguay (0.97), Pakistan (0.98), and Laos (0.99) have price levels lower than India. In spite of many of the African countries being poorer than India, only one of those listed here has a lower price level.

The final six paired columns of Table 1 show our calculations of the aggregate prices of consumption according to the three aggregation formulas and the two sources of weights. The immediate impression is that, in spite of the different weighting schemes, and different procedures, our indexes are close to the official ones. The correlation with the ICP price of consumption across the 62 countries is 0.9275 and 0.9337 for the survey and national accounts versions of the EKS-Fisher, 0.9307 and 0.9360 for the EKS-Törnqvist, and 0.9256 and 0.9346 for the weighted CPD; note that these are not correlations for the raw P3s, which would be artificially inflated by the variation in units from country to country, but the correlations of the price of consumption, whose magnitude is comparable across countries.

Table 2 explores the similarity and differences in the indexes in a more transparent, way. The top panel of the table presents distances between pairs of indexes using the root mean squared differences over countries for each pair of indexes. The first important finding is that the distances in the first row are larger than any of the others, showing that the official ICP number is further away from all of our indexes (RMSEs around 0.15 to 0.16) than any of our indexes are from one another. The ICP index and our national-accounts based indexes use the same information, but differ for two reasons. One is that our indexes are calculated in one step using a single aggregation formula, rather than different aggregation formulas by region. The second is that our indexes use only 102 of the 105 consumption basic heads in the ICP; we exclude rental (actual and imputed),

FISIM, and prostitution in order to match our National Accounts based (NAS) and survey results. As we shall see in Section 4.3, these differences have substantial effects on the calculated P3s. In terms of Table 2, recalculating the NAS based PPPs using 105 basic headings, instead of 102, reduces the MSE with the Fisher NAS index, 0.156 in Table 2, to 0.099 (not shown), with the remainder of the discrepancy coming from the different methods of calculation.

The distances between the survey and national accounts based (102 basic heads) versions of our consumption price indexes are only 0.065 (Fisher), 0.048 (Törnqvist) and 0.078 (CPD), less than half the size of the difference between our survey based indexes and the ICP national accounts based indexes. These differences are important, but smaller than the differences induced by the combination of dropping some basic heads and using the ICP method of calculation. The top panel of Table 2 also shows that the EKS-Fisher and EKS-Törnqvist indexes are typically close to one another—whether the weights come from surveys or from national accounts—and that both are somewhat further away from either of the weighted CPD indexes, a result that is consistent with the approximation ideas in Section 2 but should not be taken as an endorsement of EKS versus CPD indexes. Within a weighting scheme—national accounts or surveys—different indexes tend to be closer to one another than are the same indexes across weighting schemes. The overall conclusion is that the most important difference comes from the procedures used in the ICP versus those adapted here, as well as the exclusion of three basic heads, the second most important difference is between whether the aggregate expenditure weights come from the surveys or from the national accounts, and the least

important difference is the choice of formula, with Fisher and Törnqvist closer to one another than is either to the weighted CPD.

The second panel shows the means and standard deviations of the indexes. The standard deviations are very similar, but the ICP mean is about 3 percent lower than the others. Put differently, and in comparison with the direct calculations, the regional structure of the ICP, and other differences in calculation results in the Indian consumption price level being higher relative to the other countries listed here. The dropping of the three basic heads turns out not to be important; replacing them and recalculating the NAS-based PPPs with 105 basic heads gives the same estimates as with 102 basic heads.

The final panel of Table 2 shows a series of regressions that test for systematic differences between the national accounts and survey versions of our indexes; these help understand why the indexes differ, but will also help impute indexes for countries where we have national accounts but no survey estimates. The estimates show that survey estimates are lower in better-off countries, with the ratio falling by between one and two percent for every doubling of per capita income. Even so, the effects are barely significant. The *F*-statistics for the regional effects are typically close to significance at five percent level, but tend to be inconsistent across indexes and quite small. It is not clear whether it would be worth while using these results to estimate survey-based indexes in countries without surveys, rather than simply using the national accounts based indexes themselves.

We have looked in more detail at the reasons for the differences between the national accounts and the survey-based indexes. Since both indexes use the same parities for the

102 basic headings, differences are driven entirely by the pattern of expenditures over the parities. We have calculated, for each survey, the correlation between the (processed) survey-based estimates of the aggregate budget shares and those from the national accounts, for all categories of consumption and for the subgroup of food, drinks, tobacco and narcotics. It is not obvious what to expect of these numbers, nor how low a correlation would be a source for concern. There are a few very low numbers, even if we confine ourselves to the somewhat easier to measure food category. In an extreme case, the budget shares from the 2003 survey of Chad correlate with the national accounts numbers at only 0.090 over all goods, and only 0.023 for foods. There are a number of other correlations under 0.5. We have done some cross-checking of these numbers, and as is usually the case in comparing surveys and national accounts, the problems are not easily attributable to one side or the other.

Table 3 presents the standard errors associated with the plutocratic survey-based PPPs. We show only the EKS-Fisher and the weighted CPD; the results for the EKS-Törnqvist are similar to those for the EKS-Fisher, and indeed the estimates of the sampling standard errors are identical. We present the PPPs themselves here, rather than price of consumption; the former is the latter multiplied by the market rate of exchange of local currency to rupees. The standard errors are the standard errors of the logarithms of the PPPs, and so can be thought of as relative standard errors. They are also the standard errors for the logarithms of the prices of consumption in Table 1. There are two main points to note. First, the sampling errors are very small. Although some of the surveys have small sample sizes, the sampling standard errors for the PPP indexes are negligible. Second, the same is not true for the standard errors associated with failure of arbitrage.

Akin to the Paasche-Laspeyres spread, these standard errors measure the uncertainty associated with picking one particular index number when relative prices are not the same in different countries. These standard errors are typically in the vicinity of eight to ten percent, as opposed to a half to a tenth of one percent for the sampling standard errors. This finding of negligible standard errors from sampling, but substantial uncertainty from variations in relative prices, characterizes all of our results.

4.2 Poverty-weighted purchasing power parities, P4s

Table 4 shows the first set of poverty-weighted PPPs or P4s; these are calculated using all 50 poverty lines that we have available according to the second variant described in Section 3. Column 1 shows the Törnqvist approximation to the PPP that serves as the starting point for the further calculation, followed by the iteratively calculated Törnqvist indexes at bandwidths of 1, 0.5, and 0.1 standard deviations of the log per capita total expenditure. To average around the poverty line in (16), we use the bi-weight kernel

$$\begin{aligned} K(t) &= \frac{15}{16}(1-t^2)^2 \quad if\ |t| \le 1 \\ K(t) &= 0 \quad if\ |t| > 1 \end{aligned} \tag{21}$$

where t is the bandwidth-scaled difference between the household's per capita total expenditure and the local currency version of the international poverty line, see (16). The final two columns show the Fisher and weighted CPD P4s, both calculated using the smallest (0.1 standard deviation) bandwidth. The Törnqvist-approximation starting value is something of an outlier relative to the other indexes which are once again very similar to one another. Choosing a good bandwidth is a question of trading off bias against variance; a small bandwidth means we only use households near the poverty line, but the

result is a larger sampling variance in our estimates. Tables 5 and 6 show how this works; Table 5 lists the numbers of households at each bandwidth for the indexes in Table 4, while Table 6 lists the corresponding standard errors of the log PPPs. For example, in Table 5, we see that for a country with a large survey such as Indonesia, there are 22,760 households in the band around the poverty line when the bandwidth is 1 standard deviation, which falls to 10,415 with a bandwidth of a half, and only 1,916 with a bandwidth of 0.1. The corresponding sampling standard errors in Table 6 (multiplied by 100 compared with Table 3) rise from 0.06 to 0.08 to 0.15 of one percent so that, even with the smallest bandwidth, the sampling errors are negligible. Even for countries with much smaller sample sizes in the surveys, where the standard errors are correspondingly larger, for example Paraguay, the sampling standard errors at the smallest bandwidth are not much more than one percent.

Table 7 extends Table 2 and shows the root mean square difference, of the distances between the various indexes expressed, as before, as the price of consumption. In this table, F, T, and C stand for Fisher, Törnqvist, and CPD, respectively, while N and S stand for national accounts and surveys so that, for example, F(S) and T(N) are the plutocratic Fisher index using survey weights and the plutocratic Törnqvist index using expenditure weights from the national accounts. The indexes with numbers refer to the bandwidth, so that F1, F0.5, and F0.1 are the Fisher P4 prices of consumption calculated at bandwidths of 1, 0.5, and 0.1 of a standard deviation of the logarithm of per capita household expenditure. The first row shows, as expected, that the ICP price levels of consumption are relatively far away from the other indexes, with distances around 0.15 to 0.18. Our recalculated national accounts indexes are closer to the P4 indexes, and their survey-

based counterparts are closer still. The three national accounts P3 indexes are between 0.09 and 0.11 away from the Fisher and Törnqvist P4s, and 0.14 to 0.17 from the CPD version of the P4. The survey based P3 indexes, which use the same data as the P4s, are closer, about 0.05 to 0.07 away from the Fisher and Törnqvist and 0.09 and 0.12 for the CPD. The closed-form Törnqvist approximation that we use to start the iterations for the P4s is about as far away from the final P4s as the plutocratic survey based indexes, so these latter could just as well have been used for starting values. Once we look within the P4 indexes alone, changing the bandwidth does not move the indexes apart by much, especially within a specific index, though, as is to be expected, the adjacent bandwidths are closer than are the two extremes. Even here, the CPD P4 is not only further away from the other two indexes than they are from one another, but it also shows the largest internal changes as the bandwidth is reduced.

Table 8 examines the effects of different global poverty line procedures on the poverty-based purchasing power parity indexes. We consider two alternatives corresponding to the variants discussed in Section 3; multiplying the 50 poverty lines by two, and the Ravallion, Chen and Sangraula (2009) procedure using only 14 very poor countries.

Table 8 shows that the different assumptions do not have much effect on the poverty-weighted indexes. Replacing (a) the 50 lines with poverty weighting by (b) 14 of the 15 poorest country lines used by RCS (we have no survey data for one country) with no weighting, makes very little difference, with distances from the original consumption prices of 0.014 and 0.013 for the Fisher and Törnqvist, and of 0.036 for the CPD. Doubling the poverty lines moves the indexes somewhat further, though the distances are

only 0.050 for the Fisher, 0.048 for the Törnqvist, and 0.084 for the CPD, comparable to the distance moved by shifting from the survey based P3s to P4s. The means of the original and RCS consumption prices are close, with some increase when we double the underlying poverty lines; this presumably reflects the changing balance of global poverty between India and the rest of the world as the poverty lines are moved up, though the exact mechanism is not obvious. Once again the CPD indexes are not only further away from the Fisher and Törnqvist than they are from one another, but the CPD indexes are less internally stable, moving further when we vary the underlying poverty lines.

Table 9 looks for systematic patterns by income and region between the P4 and P3 indexes. In these regressions, the dependent variable is the logarithm of the ratio of the P4—using bandwidths of 0.1 standard deviations—to our calculated P3s using the national accounts weights. The reason for this choice is that these P3s are available for countries where there are no survey data, and are therefore the starting point for imputing P4s in the absence of survey data. None of the estimated regression coefficients are significant at conventional levels, so an argument could be made for simply using the P3 indexes. Even so, comparison with the results in Table 2, which compared the survey and national accounts based P3s, shows that the income effects here are similar, so that most of the difference between the P4s and P3s can be traced to differences between the surveys and the national accounts expenditure patterns, consistently with other evidence on the indexes.

4.3 Poverty estimates

While our main focus in this paper is the calculation of poverty-weighted purchasing power parity exchange rates, we briefly also consider the main use of those rates, which is the estimation of global poverty. As is already clear, our P4s are relatively close to the P3s from the ICP, so that the substitution of poverty weights for plutocratic national accounts weights will not, in and of itself, make a large difference to global poverty counts. This is not true of the choice of procedure for calculating the global poverty line, including the effects of our PPPs on this calculation.

Table 10 presents a sample of our poverty estimates for the world and for its main regions, with different PPPs and different procedures for calculating the global line. The first set of numbers reproduces the Bank's poverty counts for 2005, World Bank (2008b). Their poverty line is $38 per person per month (first row), calculated as the unweighted average of the PPP value of the local lines of 15 of the world's poorest countries (second row) with the conversion done using the P3s (third row) from the ICP (fourth row.) These parameters give a global poverty total of 1.32 billion, with the distribution over regions as shown. In the next three columns, we switch to P4s, and show the three different aggregation formulas, CPD, Fisher, and Törnqvist. Because we lack household survey data for Guinea-Bissau, which is one of the 15 countries, we work with the remaining 14. This exclusion makes almost no difference, and we can reproduce the first column very closely using the 14 countries and the PPPs from the ICP (not shown here). When we switch to P4s, the global count varies from 1.13 billion using the Törnqvist to 1.21 using the CPD. (If we were to prefer the EKS indexes to the CPD, as was argued above, the range is only 1.13 to 1.16.) The reduction in global poverty from the Bank numbers, from 1.32 billion, comes primarily from our treatment of housing rental in the 14 index

countries where several of the ICP parity estimates are incorrect. (The ICP treatment of housing is entirely appropriate for their main purpose, which is the estimation of GDP, but is not appropriate for poverty calculations, see Deaton and Heston, 2009, and Deaton, 2010 for further discussion.) Our treatment of rentals, which assumes that the parity for rentals is the same as for consumption as a whole, raises the P3s and P4s for several of the index countries, which lowers their poverty lines in international currency, so that our poverty lines—shown here in international rupees given that we cannot include the US in our P4s—are lower than the Bank's, and a lower line gives a lower poverty count.

The final three columns in Table 10 show our preferred poverty estimates, in the sense that they are most closely comparable with the dollar a day counts published by the Bank prior to the 2005 round of the ICP. Here we use poverty lines, not just from the 14 poorest countries, but from the 50 countries whose poverty lines are included in RCS. Because these countries differ in levels of development and poverty rates, we use our P4s to convert their poverty lines, and then take a weighted average using as weights the numbers of poor people in each country. These calculations are done simultaneously with the calculation of the P4s, so that the international rupee value of the local poverty lines, the local poverty counts, and the P4s are all mutually consistent once the calculations are completed. These global poverty lines are sharply lower, not only than the Bank's line, but also than our own P4 lines using the 14 countries. The biggest contribution to this difference is the inclusion of India in the 50 countries. India has a very low poverty line relative to its level of GDP per capita, and makes a large contribution to global poverty, so its inclusion in the 50 country calculation brings down the global line and the global count. After the 2005 round of the ICP, the Bank recalculated its global line with the new

P3s, but also excluding India, which is the basis for our arguing that the final three columns in Table 10 are most closely comparable to the original procedures. As was the case with the 14-country based P4s, there is little difference in counts according to the aggregation formula used for the P4s.

7. Summary and conclusions

Our aim in this paper is to show how to calculate purchasing power parity exchange rates that reflect the consumption patterns of poor people around the world, poverty-based PPPs, or P4s, rather than the familiar P3s. P4s, unlike P3s, require household survey information, but there are currently enough household surveys to cover the vast majority of the world's poor population. P4s, unlike P3s, need to be calculated simultaneously with the global poverty line, because the price indexes depend on the line and the line depends on the price indexes. We have shown how the fixed point can be calculated explicitly in a special case, and developed an iterative procedure that works more generally. We have developed formulas for standard errors of our estimates in order to address the concern that some of the household surveys have small samples, so that the estimates might be too noisy for use. In practice, the standard errors from sampling are very small, negligibly so relative to the more general uncertainty associated with the choice of index number formula. The design and detail of household surveys vary widely across the world, and many compromises and assumptions have to be made to adapt the survey data to match the prices from the International Comparison Project. In the end, we believe our procedures are unlikely to be a source of much error in our final estimates.

In the end, poverty-weighted purchasing power parity exchange rates look very much like the regular purchasing power parity exchange rates that use weights from the national accounts, certainly when we confine ourselves to comparisons that do not involve the rich countries of the world. Although it is true that poor people have different consumption patterns from the patterns in the national accounts, the reweighting is similar in different countries, so that the price indexes between each pair do not usually change by much. There are, of course, exceptions, but the weighting differences between P4s and P3s are probably not of great importance for estimating global poverty.

A larger source of difference between the P3s and P4s is data inconsistency between household surveys and national accounts, so that the consumption pattern in one is often different from the consumption pattern in the other, even when we use both to estimate aggregate consumption. Some of this comes from difference in definition and coverage—FISIM and owner-occupied rental equivalence are not collected in surveys, nor (usually) are expenditures on narcotics or prostitution. Perhaps more important are measurement errors in either the surveys or the national accounts or both. Yet even the differences in these weights do not generate large differences between P3s and P4s.

There are a number of important issues that we do not address. Leading among these is the fact that we make no attempt to use separate *prices* for the poor. Instead, we confine ourselves to reweighting the same prices to match the expenditure patterns of households near the global poverty line. This is of particular concern in the 2005 round of the ICP, which used very detailed specifications of the goods to be priced, so that, in some countries and for some goods, the ICP prices may reflect, not typical local prices, but the prices of international goods in a few high-priced stores in the capital patronized

by the rich and by expatriates. The Asian Development Bank (2008) has undertaken experimental work to identify the prices paid by the poor, by collecting prices in shops and markets thought to be patronized by the poor, and by specifying varieties of goods that are typically purchased by the poor. One potential weakness of these procedures is that it is unclear exactly what and where the poor buy, and the ADB's specifications were set by groups of experts. Perhaps a better source of such information is to use the unit values in household surveys, which have the advantage of relating to actual purchases by poor people. The corresponding disadvantage is that there is no obvious way of specifying quality, or of controlling for quality variation across poor and non-poor. A useful project would be to compare unit values with the prices collected for the 2005 ICP.

Our work also raises a number of issues that are relevant both for future work on the ICP and on household surveys. For the former, it is clear that, in some respects, the demands of national accounting and of poverty work are different. For example, for poverty work we need prices paid by consumers, not prices paid by governments on behalf of consumers, a distinction that is particularly troubling in the case of health related goods, such as pharmaceuticals. It is also the case that when direct measurements break down or are difficult, the supplementary imputations that are suitable for estimating national accounts are sometimes different than those that would make most sense for estimating poverty.

On household surveys, our plea is mostly for greater harmonization across countries. We realize that surveys are used for different purposes in different countries, and that a survey that works in one country may be useless in another. Nevertheless, greater

standardization is certainly possible in some cases, not only in data collection, but in the reporting and documentation of survey design.

Although we suspect that it is not of leading importance for the estimates presented here, we also want to flag the issue of quality adjustment. How to deal with quality is perhaps the leading unsolved issue in price index construction, both domestically—see for example Mackie and Schultze (2002) for the US—and internationally in the ICP. The ICP has become progressively more detailed in comparing like with like across countries, on the reasonable suspicion that price levels in poor countries were being understated by comparing lower quality goods in poor countries with higher quality goods in richer countries. The use of more precise specifications has raised price levels in poor countries in more recent rounds. It is certainly true that the latest ICP does a better job of the quality comparisons, though perhaps at the price in some cases of comparing representative goods in a rich country with the same goods in a poor country but which are not representative of consumption patterns there. The 2005 ICP attempted to make a representativity correction to deal with this but, for a number of reasons, the correction was not successful for poor countries. This is an active area of future research for the ICP itself, and is likely to have repercussions for poverty work in the future.

8. List of references

Asian Development Bank, 2008, *Research study on poverty-specific purchasing power parities for selected countries in Asia and the Pacific*, Manila, March.

Chen, Shaohua, and Martin Ravallion, 2008, "The developing world is poorer than we thought, but no less successful in the fight against poverty," Policy Research Working Paper 4703, World Bank, August.

Deaton, Angus, 2003, "Prices and poverty in India, 1987–2000," *Economic and Political Weekly*, January 25th, 362–8.

Deaton, Angus, 2005, "Measuring poverty in a growing world (or measuring growth in a poor world," *Review of Economics and Statistics*, 87: 1, 1–19.

Deaton, Angus, and Olivier Dupriez, 2009, "Global poverty and global price indexes," World Bank and Princeton, processed. http://www.princeton.edu/~deaton/downloads/Global_Poverty_and_Global_Price_Indexes.pdf

Deaton, Angus, 2010, "Quality, inequality, and the measurement of global poverty," *American Economic Review*, March. (forthcoming.)

Deaton, Angus, and John Muellbauer, 1980, *Economics and Consumer Behavior*, New York. Cambridge University Press.

Deaton, Angus, and Alan Heston, 2008, "Understanding PPPs and PPP-based national accounts," NBER Working Paper No. 14499, November.

Deaton, Angus and Sam Schulhofer-Wohl, 2009, "International comparisons of well-being," in preparation.

Diewert, W. Erwin, 1976, "Exact and superlative indexes," *Journal of Econometrics*, 4:2, 115–45.

Diewert, W. Erwin, 2008, "New methodology for linking the regions," DP 08-07, Department of Economics, University of British Columbia, processed, July.

Gini, Corrado, 1924, "Quelques considérations au sujet de la construction des nombres indices des prix et des questions analogues," *Metron*, Vol. IV N.1, 3-162, Roma.

Hill, Peter, 2007a, "The Ring Program: Linking the Regions", Chapter 13 in *ICP 2003-2006 Handbook*, Washington D.C.: The World Bank. http://siteresources.worldbank.org/ICPINT/Resources/Ch13_Ring_Feb07.doc

Hill, Peter, 2007b, "Ring Comparison—Linking Within-Region PPPs Using Between Region PPPs", Chapter 14 in *ICP 2003-2006 Handbook*, Washington D.C.: The World Bank. http://siteresources.worldbank.org/ICPINT/Resources/ch14_Linking_Apr_06.doc

Kulshreshtha, A. C., and Aloke Kar, 2005, "Estimates of food consumption expenditure from household surveys and national accounts," in Deaton, Angus and Valerie Kozel, *The great Indian poverty debate*, Delhi, Macmillan, 102–18.

Prais, Sigbert J., 1959, "Whose cost of living?" *Review of Economic Studies*, 26, 126–34.

Rao, D.S. Prasada, 1990, "A System of Log-Change Index Numbers for Multilateral Comparisons", pp. 127-139 in *Comparisons of Prices and Real Products in Latin America*, J. Salazar-Carillo and D.S. Prasada Rao (eds.), New York: Elsevier Science Publishers.

Rao, D.S. Prasada, 2005, "On the Equivalence of the Generalized Country-Product-Dummy (CPD) Method and the Rao System for Multilateral Comparisons", *Review of Income and Wealth*, 51:4, 571–80.

Ravallion, Martin, Shaohua Chen, and Prem Sangraula, 2009, "Dollar a day revisited," *World Bank Economic Review*. 23:2, 163–84.

Mackie, Christopher D and Charles L. Schultze, 2002, *At what price? Conceptualizing and measuring cost of living and price indexes*, Washington, National Academies Press.

Selvanathan, E. A. and D. S. Prasada Rao, 1994, *Index Numbers: A Stochastic Approach*, Ann Arbor: The University of Michigan Press.

Summers, Robert, 1973, "International comparisons with incomplete data," *Review of Income and Wealth*, 29:1, 1–16.

Triplett, Jack E., 1997, "Measuring consumption: the post-1973 slowdown and the research issues," *Federal Reserve Bank of St Louis Review*, May/June, 9–42.

World Bank, 1990, *World Development Report: Poverty*, New York, Oxford University Press.

World Bank, 2008a, *Global purchasing power parities and real expenditures: 2005 international comparison program*, Washington, DC.

World Bank, 2008b, *Poverty data: a supplement to the World Development Indicators 2008*, Washington, DC.

Table 1: Consumption prices using national aggregates as weights

	Pc ICP	Pc Fisher		Pc Törnqvist		Pc-CPD(W)	
		NAS	Survey	NAS	Survey	NAS	Survey
India	1	1	1	1	1	1	1
Bangladesh	1.120	1.091	1.077	1.079	1.063	1.098	1.073
Bhutan	1.183	1.158	1.139	1.135	1.128	1.126	1.142
Cambodia	1.116	1.111	1.175	1.092	1.147	1.057	1.135
China	1.411	1.404	1.354	1.410	1.389	1.399	1.361
Fiji	2.589	2.222	2.124	2.162	2.106	2.184	2.079
Indonesia	1.221	1.185	1.184	1.163	1.169	1.143	1.168
Lao PDR	0.993	1.043	1.090	1.048	1.076	1.033	1.123
Malaysia	1.577	1.497	1.440	1.471	1.439	1.416	1.379
Maldives	2.150	1.716	1.721	1.708	1.702	1.668	1.613
Mongolia	1.225	1.217	1.234	1.204	1.216	1.166	1.172
Nepal	1.048	0.989	1.003	0.976	0.999	0.950	0.999
Pakistan	0.984	1.038	1.071	1.029	1.055	1.005	1.052
Philippines	1.241	1.238	1.249	1.221	1.238	1.194	1.199
Sri Lanka	1.126	1.178	1.150	1.157	1.142	1.128	1.106
Thailand	1.227	1.306	1.299	1.268	1.273	1.219	1.232
Vietnam	1.055	1.031	1.058	1.044	1.069	1.028	1.048
Argentina	1.318	1.383	1.347	1.374	1.359	1.363	1.326
Bolivia	0.900	1.020	1.056	1.007	1.043	0.955	1.013
Brazil	1.828	1.992	1.951	1.912	1.888	1.956	1.917
Colombia	1.452	1.676	1.693	1.642	1.644	1.619	1.595
Paraguay	0.974	1.094	1.074	1.083	1.074	1.051	1.030
Peru	1.416	1.670	1.621	1.642	1.571	1.677	1.540
Armenia	1.212	1.146	1.164	1.142	1.143	1.140	1.124
Azerbaijan	1.039	0.961	0.883	0.968	0.933	0.987	0.918
Kazakhstan	1.382	1.070	1.060	1.100	1.068	1.122	1.068
Kyrgyz Rep.	0.896	0.789	0.807	0.823	0.837	0.822	0.856
Tajikistan	0.840	0.613	0.775	0.783	0.821	0.755	0.844
Yemen	1.345	1.201	1.150	1.166	1.156	1.139	1.150

Notes: The first column is from the ICP Final Report, and is the PPP for individual consumption expenditures by households divided by the foreign exchange rate, the "price of consumption" with India as base. The second, third, and fourth columns report are prices of consumption using the parities for 102 basic heads, but using estimates of aggregate weights first from the national accounts, then from the household surveys. The first column and the first column of each pair differ only in the aggregation formulas, the ring structure, and the merging of regional parities for the basic headings of consumption.

Table 1, continued.

	Pc ICP	Pc Fisher		Pc Törnqvist		Pc CPD (W)	
		NAS	Survey	NAS	Survey	NAS	Survey
Benin	1.475	1.545	1.448	1.544	1.490	1.576	1.499
Burkina Faso	1.299	1.417	1.382	1.389	1.379	1.388	1.376
Burundi	1.168	1.283	1.214	1.298	1.212	1.301	1.163
Cameroon	1.578	1.690	1.681	1.674	1.686	1.665	1.655
Cape Verde	2.493	2.402	2.295	2.383	2.286	2.382	2.264
Chad	1.755	1.995	1.882	1.944	1.847	2.082	1.849
Congo DR	1.886	1.975	1.989	1.934	1.961	1.976	2.010
Congo PR	2.013	2.122	2.072	2.111	2.072	2.122	2.083
Côte d'Ivoire	1.746	1.850	1.828	1.837	1.846	1.859	1.850
Djibouti	1.715	1.950	2.051	1.935	2.025	1.796	1.985
Ethiopia	0.897	1.068	1.039	1.035	1.016	0.982	0.978
Gabon	2.378	2.505	2.469	2.507	2.483	2.565	2.525
Gambia	1.023	1.224	1.314	1.232	1.296	1.147	1.247
Ghana	1.394	1.593	1.540	1.577	1.540	1.572	1.516
Guinea	1.148	1.260	1.254	1.272	1.270	1.310	1.328
Kenya	1.223	1.380	1.340	1.370	1.335	1.377	1.326
Lesotho	1.523	1.671	1.726	1.712	1.721	1.650	1.677
Madagascar	1.066	1.111	1.153	1.132	1.159	1.171	1.211
Malawi	1.359	1.572	1.462	1.577	1.501	1.559	1.482
Mali	1.552	1.663	1.585	1.641	1.590	1.654	1.601
Mauritania	1.341	1.569	1.530	1.534	1.507	1.521	1.469
Morocco	1.756	1.929	1.777	1.897	1.800	1.901	1.772
Mozambique	1.409	1.658	1.471	1.616	1.477	1.578	1.395
Niger	1.433	1.602	1.575	1.579	1.570	1.567	1.575
Nigeria	1.692	1.836	1.826	1.827	1.824	1.874	1.848
Rwanda	1.200	1.287	1.352	1.284	1.375	1.211	1.331
Senegal	1.598	1.768	1.742	1.751	1.727	1.758	1.696
Sierra Leone	1.361	1.597	1.571	1.593	1.576	1.539	1.510
South Africa	2.032	2.172	2.034	2.129	2.013	2.168	2.016
Swaziland	1.657	1.815	1.709	1.816	1.726	1.761	1.590
Tanzania	1.218	1.304	1.267	1.269	1.248	1.284	1.257
Togo	1.513	1.644	1.595	1.631	1.605	1.681	1.618
Uganda	1.182	1.240	1.172	1.257	1.205	1.230	1.154

Table 2: Survey based and NAS based estimates of the price of aggregate consumption

	ICP	Fisher (N)	Fisher (S)	Törnqvist (N)	Törnqvist (S)	CPD (N)	CPD (S)
				Root mean square distance			
ICP	0	0.156	0.150	0.147	0.146	0.149	0.148
Fisher (N)		0	0.065	0.033	0.068	0.050	0.088
Fisher (S)			0	0.054	0.023	0.078	0.047
Törnqvist (N)				0	0.048	0.042	0.067
Törnqvist (S)					0	0.066	0.070
CPD (N)						0	0.078
CPD (S)							0
				Summary statistics			
Mean	1.402	1.463	1.440	1.453	1.437	1.445	1.421
Standard dev.	0.389	0.404	0.377	0.390	0.372	0.404	0.373
				Regressions of log of ratio of Survey to National Accounts basis			
		Estimate	*t*-value	Estimate	*t*-value	Estimate	*t*-value
ln *y*		−0.0170	(2.1)	−0.0107	(2.1)	−0.0200	(2.3)
Asia		0.0055	(0.1)	0.0077	(0.3)	0.0143	(0.3)
Africa		−0.0334	(0.7)	−0.0221	(0.7)	−0.0345	(0.7)
Latin America		0.0086	(0.2)	0.0041	(0.1)	0.0019	(0.0)
Central Asia		0.0283	(0.6)	0.0020	(0.0)	0.0011	(0.2)
constant		0.1313	(1.5)	0.0825	(1.7)	0.1542	(1.7)
F-regions (*p*)		2.69	0.041	2.53	0.051	2.97	0.056

Notes: The top panel shows the root mean squared difference between pair of consumption price indexes over the 62 countries. The country price indexes are those shown in Table 1. Means and standard deviations in the second panel refer to the same indexes. The final panel shows regressions of the log of the ratio of the survey-based to national accounts based estimates on the log of per capita GDP in PPP $ (from the 2008 World Development Indicators) and dummies for the ICP regions. For these regressions, India is treated as a region, and is the base country, so that Asia refers to non-Indian Asia.

Table 3: PPPs for consumption using national aggregates from surveys, and the standard errors of their logarithms

	Pc Fisher			Pc-CPD(W)		
	PPP	se(1)	se(2)	PPP	se(1)	se(2)
India	1.000	---	---	1.000	---	---
Bangladesh	1.571	0.0010	0.0836	1.565	0.0040	0.1048
Bhutan	1.139	0.0012	0.0693	1.142	0.0025	0.0828
Cambodia	109.1	0.0007	0.1040	105.4	0.0027	0.1308
China	0.251	0.0004	0.0975	0.253	0.0029	0.1293
Fiji	0.081	0.0011	0.0815	0.080	0.0038	0.0967
Indonesia	260.6	0.0004	0.0757	257.0	0.0026	0.0940
Lao PDR	263.2	0.0040	0.1000	271.4	0.0027	0.1370
Malaysia	0.124	0.0052	0.0862	0.118	0.0035	0.1128
Maldives	0.499	0.0062	0.0954	0.468	0.0038	0.1219
Mongolia	33.73	0.0007	0.0851	32.02	0.0031	0.1039
Nepal	1.622	0.0014	0.0848	1.616	0.0090	0.1046
Pakistan	1.446	0.0005	0.0799	1.420	0.0039	0.0941
Philippines	1.560	0.0005	0.0858	1.498	0.0040	0.1040
Sri Lanka	2.621	0.0006	0.0861	2.521	0.0032	0.1051
Thailand	1.185	0.0005	0.0765	1.124	0.0028	0.0877
Vietnam	380.6	0.0010	0.0860	376.9	0.0033	0.1104
Argentina	0.089	0.0008	0.0813	0.087	0.0045	0.0982
Bolivia	0.193	0.0015	0.0790	0.185	0.0048	0.0946
Brazil	0.107	0.0012	0.0986	0.106	0.0035	0.1196
Colombia	89.07	0.0011	0.0795	83.93	0.0034	0.0938
Paraguay	150.5	0.0017	0.0830	144.3	0.0034	0.1026
Peru	0.121	0.0010	0.0798	0.115	0.0038	0.0906
Armenia	12.08	0.0025	0.0791	11.66	0.0039	0.0894
Azerbaijan	94.62	0.0043	0.0950	98.37	0.0039	0.1164
Kazakhstan	3.195	0.0006	0.0809	3.219	0.0106	0.0921
Kyrgyz Rep.	0.751	0.0041	0.0969	0.796	0.0049	0.1076
Tajikistan	0.055	0.0026	0.0974	0.060	0.0052	0.1061
Yemen	4.993	0.0017	0.0868	4.991	0.0033	0.1035

Notes: Pc is the aggregate (plutocratic) consumption PPP expressed in local currency per Indian rupee. The Törnqvist is not shown because the results are similar to those for the Fisher index. The second and third columns of each set show (a) the standard errors associated with sampling from the household surveys and (b) the standard errors associated with the failure of arbitrage. Standard errors are standard errors of the logarithms of the PPPs shown in the first column. Standard errors for India and China are not shown; the former is the base country, while for China we are using synthetic data that matches the published tables.

Table 3, continued: PPPs for consumption using national aggregates from surveys, and their standard errors

	Pc Fisher			Pc-CPD(W)		
	PPP	se(1)	se(2)	PPP	se(1)	se(2)
Benin	17.32`	0.0014	0.0966	17.93	0.0057	0.1323
Burkina Faso	16.53	0.0011	0.0746	16.45	0.0032	0.0906
Burundi	29.78	0.0022	0.1077	28.52	0.0047	0.1544
Cameroon	20.11	0.0014	0.0715	19.79	0.0028	0.0855
Cape Verde	4.613	0.0022	0.0893	4.551	0.0031	0.1051
Chad	22.52	0.0012	0.0742	22.12	0.0023	0.0884
Congo DR	21.37	0.0008	0.0706	21.60	0.0033	0.0867
Congo PR	24.78	0.0012	0.0755	24.92	0.0027	0.0883
Côte d'Ivoire	21.86	0.0018	0.0741	22.12	0.0034	0.0906
Djibouti	8.267	0.0010	0.0774	7.999	0.0041	0.0970
Ethiopia	0.204	0.0013	0.0846	0.192	0.0055	0.0970
Gabon	29.54	0.0009	0.0805	30.20	0.0030	0.0942
Gambia	0.852	0.0025	0.0800	0.808	0.0030	0.0935
Ghana	316.8	0.0009	0.0751	312.0	0.0069	0.0866
Guinea	103.7	0.0019	0.0975	109.8	0.0028	0.1237
Kenya	2.295	0.0010	0.0703	2.272	0.0026	0.0847
Lesotho	0.249	0.0019	0.0752	0.242	0.0032	0.0900
Madagascar	52.44	0.0023	0.0817	55.06	0.0039	0.0984
Malawi	3.927	0.0031	0.1121	3.980	0.0037	0.1549
Mali	18.96	0.0008	0.0710	19.15	0.0036	0.0859
Mauritania	9.190	0.0009	0.0751	8.823	0.0047	0.0900
Morocco	0.357	0.0008	0.0923	0.356	0.0033	0.1095
Mozambique	777.9	0.0030	0.0989	737.5	0.0031	0.1317
Niger	18.84	0.0011	0.0723	18.83	0.0024	0.0883
Nigeria	5.435	0.0011	0.0861	5.500	0.0029	0.1009
Rwanda	17.10	0.0021	0.0971	16.83	0.0031	0.1273
Senegal	20.83	0.0006	0.0700	20.28	0.0031	0.0843
Sierra Leone	103.3	0.0025	0.0848	99.26	0.0077	0.0989
South Africa	0.293	0.0014	0.0832	0.291	0.0030	0.1004
Swaziland	0.246	0.0040	0.0831	0.229	0.0027	0.1068
Tanzania	32.15	0.0013	0.0743	31.91	0.0046	0.0887
Togo	19.08	0.0009	0.0775	19.35	0.0029	0.0912
Uganda	47.33	0.0019	0.1105	46.58	0.0033	0.1536

Table 4: Poverty-weighted PPPs at various bandwidths

	Törnqvist Indexes				Fisher	CPD(W)
Bandwidth	Approx.	1.0	0.5	0.1	0.1	0.1
India	1.000	1.000	1.000	1.000	1.000	1.000
Bangladesh	1.479	1.501	1.496	1.494	1.517	1.510
Bhutan	1.114	1.089	1.086	1.086	1.098	1.081
Cambodia	102.9	103.0	102.5	102.3	104.0	100.2
China	0.252	0.253	0.252	0.252	0.246	0.241
Fiji	0.082	0.080	0.080	0.080	0.081	0.077
Indonesia	259.3	252.5	251.5	251.0	255.3	245.5
Lao	260.6	251.8	251.3	252.7	256.1	260.3
Malaysia	0.128	0.124	0.124	0.123	0.125	0.117
Maldives	0.532	0.506	0.501	0.491	0.505	0.484
Mongolia	33.84	32.92	32.83	32.74	33.23	30.65
Nepal	1.487	1.535	1.532	1.531	1.539	1.514
Pakistan	1.490	1.438	1.440	1.439	1.457	1.396
Philippines	1.522	1.482	1.476	1.473	1.486	1.382
Sri Lanka	2.554	2.521	2.514	2.509	2.526	2.346
Thailand	1.183	1.121	1.113	1.120	1.156	0.963
Vietnam	359.0	357.8	355.4	354.3	354.7	336.3
Argentina	0.083	0.081	0.081	0.081	0.080	0.073
Bolivia	0.192	0.183	0.183	0.183	0.186	0.175
Brazil	0.100	0.101	0.101	0.102	0.103	0.106
Colombia	93.99	87.81	87.78	88.47	89.99	86.33
Paraguay	147.8	144.9	144.1	145.2	145.4	138.0
Peru	0.122	0.117	0.117	0.118	0.120	0.114
Armenia	12.29	11.56	11.51	11.51	11.68	10.97
Azerbaijan	96.61	95.28	95.34	96.47	89.41	90.59
Kazakhstan	2.999	2.998	2.998	2.998	3.006	2.890
Kyrgystan	0.799	0.755	0.744	0.741	0.740	0.715
Tajikistan	0.060	0.055	0.056	0.056	0.048	0.054
Yemen	4.885	4.781	4.750	4.795	4.631	4.494

Notes: Authors calculations using formulas described in the text. These are based on 50 local poverty lines, and use 102 basic heads. The global poverty line is calculated by weighting each country's poverty line in international wrupees by the estimated number of people below the line in that country.

Table 4, continued, poverty-weighted PPPs at various bandwidths

	Törnqvist Indexes				Fisher	CPD(W)
Bandwidth	Approx.	1.0	0.5	0.1	0.1	0.1
Benin	18.26	17.70	17.68	17.64	17.14	16.87
Burkina Faso	16.12	15.97	15.93	15.91	15.90	15.40
Burundi	30.07	29.19	29.17	29.10	29.20	26.63
Cameroon	20.08	19.69	19.68	19.65	19.59	18.62
Cape Verde	4.308	4.297	4.273	4.303	4.354	4.067
Chad	23.17	22.12	22.11	22.10	22.48	21.46
Congo DR	21.30	20.93	20.91	20.88	21.10	20.81
Congo PR	26.19	24.68	24.66	24.67	24.52	23.89
Côte d'Ivoire	22.20	21.62	21.58	21.52	21.47	21.16
Djibouti	8.169	7.999	7.974	8.024	8.122	7.913
Ethiopia	0.200	0.194	0.194	0.193	0.197	0.178
Gabon	29.91	29.47	29.56	29.77	29.57	30.67
Gambia	0.912	0.855	0.855	0.853	0.859	0.785
Ghana	350.4	322.8	323.0	322.2	321.8	302.3
Guinea	111.3	105.8	105.8	105.8	104.2	109.0
Kenya	2.287	2.242	2.239	2.237	2.256	2.141
Lesotho	0.253	0.242	0.242	0.241	0.241	0.226
Madagascar	54.38	53.06	52.98	52.97	52.33	56.57
Malawi	3.993	3.909	3.903	3.887	3.782	3.622
Mali	19.29	18.73	18.70	18.71	18.60	18.38
Mauritania	9.466	8.942	8.919	8.875	9.066	8.415
Morocco	0.351	0.336	0.330	0.326	0.324	0.292
Mozambique	707.9	718.9	715.7	714.7	709.6	637.6
Niger	19.02	18.51	18.49	18.48	18.56	17.99
Nigeria	6.217	5.604	5.610	5.621	5.604	5.352
Rwanda	17.32	16.93	17.00	16.99	16.64	15.73
Senegal	21.13	20.28	20.24	20.24	20.45	19 45
Sierra Leone	107.8	103.5	103.5	103.2	102.9	96.47
South Africa	0.265	0.265	0.264	0.262	0.266	0.246
Swaziland	0.257	0.247	0.247	0.249	0.248	0.225
Tanzania	32.00	31.22	31.17	31.15	31.39	30.46
Togo	19.89	19.16	19.14	19.15	19.04	18.70
Uganda	46.74	46.15	46.04	45.76	44.34	40.90

Table 5: Numbers of observations within the bandwidth around the poverty lines (first column is total number of households in the survey)

	Sample size	T(1.0)	T(0.5)	T(0.1)	F(0.1)	CPD(0.1)
India	124644	78724	45623	9670	9761	10003
Bangladesh	7448	5595	3049	616	631	638
Bhutan	4007	1047	469	84	82	81
Cambodia	14984	7014	3392	641	683	650
*China**	*2000*	*721*	*363*	*74*	*71*	*71*
Indonesia	64422	22760	10415	1916	2098	1918
Fiji	5244	1761	807	158	158	149
Lao	8071	5589	3197	658	678	686
Malaysia	14084	363	76	11	14	8
Maldives	2728	157	42	11	11	7
Mongolia	11162	4112	1913	339	371	334
Nepal	3912	2329	1349	301	305	305
Pakistan	15839	6993	3198	573	613	547
Philippines	42094	17839	8998	1814	1882	1673
Sri Lanka	16924	4484	1785	342	360	258
Thailand	34785	414	80	8	13	5
Vietnam	9189	4224	1938	345	353	340
Argentina	27245	2304	798	135	136	109
Bolivia	5732	1125	415	77	77	72
Brazil	48466	8446	3138	568	593	635
Colombia	22949	2357	880	166	169	163
Paraguay	2682	580	260	51	47	52
Peru	18911	3464	1227	219	217	214
Armenia	6816	873	322	62	63	60
Azerbaijan	7820	1038	338	64	43	51
Kazakhstan	11986	128	44	6	6	6
Kyrgystan	1081	210	81	14	16	14
Tajikistan	4160	768	290	51	23	52
Yemen	13136	1327	460	67	71	73

* A synthetic dataset was used for China (see Appendix).

Table 5, continued: Numbers of observations within the bandwidth around the poverty lines (first column is total number of households in the survey)

	Sample Size	T(1.0)	T(0.5)	T(0.1)	F(0.1)	CPD(0.1)
Benin	5350	3552	2008	422	427	430
Burkina Faso	8494	5795	3330	674	677	685
Burundi	6668	3807	2124	444	436	463
Cameroon	10992	5111	2603	522	524	482
Cape Verde	4584	1967	965	186	197	173
Chad	6697	4279	2318	445	469	448
Congo DR	11959	6626	3508	713	709	714
Congo PR	5002	2742	1389	284	284	276
Côte d'Ivoire	10800	5473	2769	562	564	567
Djibouti	2380	794	344	45	49	49
Ethiopia	16672	7966	4206	898	956	697
Gabon	6379	1070	424	74	74	93
Gambia	2238	1326	737	167	171	137
Ghana	8687	4513	2335	443	442	442
Guinea	7095	4901	2755	571	568	569
Kenya	13154	8055	4534	942	966	932
Lesotho	5992	3532	1876	404	404	418
Madagascar	5078	996	391	56	60	82
Malawi	11280	7428	4048	838	855	889
Mali	4494	3065	1843	401	400	406
Mauritania	9385	2991	1335	245	279	219
Morocco	14243	5508	1085	96	93	70
Mozambique	8700	5931	3400	679	668	698
Niger	6689	4419	2438	528	521	532
Nigeria	19158	13019	7350	1572	1565	1574
Rwanda	6900	3326	1496	266	268	297
Senegal	6594	4095	2266	483	490	464
Sierra Leone	3719	2717	1574	352	353	353
South Africa	26215	10039	4772	948	959	913
Swaziland	3794	2907	1739	343	344	385
Tanzania	22178	13996	7670	1601	1587	1604
Togo	7500	5218	3011	616	616	616
Uganda	9711	6295	3641	755	755	737

Table 6. Estimates of standard errors of log P4s from sampling, percentages

	T(1.0)	T(0.5)	T(0.1)	F(0.1)	CPD(0.1)
India	---	0.00	0.00	0.00	0.00
Bangladesh	0.07	0.09	0.15	0.18	0.32
Bhutan	0.15	0.17	0.33	0.54	0.69
Cambodia	0.15	0.18	0.30	0.28	0.61
China	*0.05*	*0.06*	*0.13*	*0.13*	*0.30*
Indonesia	0.06	0.08	0.15	0.13	0.29
Fiji	0.16	0.24	0.58	0.87	1.25
Lao PDR	0.17	0.19	0.32	0.27	0.65
Malaysia	0.25	0.62	0.64	0.87	2.76
Maldives	0.52	0.83	1.59	1.49	3.24
Mongolia	0.18	0.21	0.38	0.30	0.79
Nepal	0.14	0.16	0.25	0.23	0.51
Pakistan	0.10	0.13	0.22	0.19	0.49
Philippines	0.09	0.12	0.20	0.20	0.37
Sri Lanka	0.10	0.13	0.26	0.25	0.62
Thailand	0.65	1.02	2.10	0.54	0.78
Vietnam	0.11	0.15	0.29	0.27	0.61
Argentina	0.19	0.32	1.09	1.13	1.06
Bolivia	0.24	0.29	0.76	0.74	1.27
Brazil	0.24	0.36	0.83	0.66	1.18
Colombia	0.19	0.31	0.65	0.61	1.25
Paraguay	0.36	0.48	1.28	1.06	2.01
Peru	0.20	0.29	0.63	0.45	1.33
Armenia	0.16	0.23	0.47	0.62	0.92
Azerbaijan	0.33	0.52	0.92	3.11	2.95
Kazakhstan	0.37	0.66	0.45	0.34	1.27
Kyrgystan	0.57	0.83	1.56	1.39	2.28
Tajikistan	0.28	0.46	1.42	0.65	2.28
Yemen	0.52	0.76	2.05	0.90	2.24

Note: The figures shown have been multiplied by 100, and are already standard errors of logs. Hence, for example, the estimated standard error of the log of the Törnqvist P4 for the Maldives with bandwidth 1 is 0.0052, or a little over half of one percent. For Armenia, Azerbaijan, Fiji, Ghana, Kazakhstan, Tajikistan, Kyrgyzstan, and Morocco, we do not have information on the survey design and have assumed that the surveys are unstratified simple random samples, so that the standard errors shown are almost certainly too small. A synthetic dataset was used for China (see Appendix).

Table 6, continued. Estimates of standard errors of log P4s from sampling, percentages

	T(1.0)	T(0.5)	T(0.1)	F(0.1)	CPD(0.1)
Benin	0.16	0.18	0.33	0.36	0.55
Burkina Faso	0.09	0.11	0.22	0.24	0.46
Burundi	0.24	0.27	0.46	0.44	1.04
Cameroon	0.25	0.28	0.41	0.53	0.74
Cape Verde	0.31	0.40	0.56	0.62	1.18
Chad	0.10	0.12	0.24	0.27	0.46
Congo DR	0.12	0.16	0.30	0.21	0.51
Congo PR	0.13	0.18	0.30	0.32	0.61
Côte d'Ivoire	0.12	0.15	0.28	0.34	0.53
Djibouti	0.19	0.29	0.53	0.68	1.02
Ethiopia	0.13	0.15	0.26	0.26	0.55
Gabon	0.20	0.30	0.68	0.70	1.15
Gambia	0.32	0.37	0.62	0.63	1.38
Ghana	0.08	0.11	0.23	0.26	0.47
Guinea	0.21	0.26	0.47	0.51	0.86
Kenya	0.08	0.09	0.17	0.22	0.34
Lesotho	0.14	0.18	0.33	0.41	0.62
Madagascar	0.20	0.26	0.54	0.57	1.12
Malawi	0.14	0.17	0.34	0.41	0.59
Mali	0.09	0.12	0.25	0.29	0.47
Mauritania	0.15	0.19	0.35	0.35	0.68
Morocco	0.13	0.26	0.79	0.87	1.68
Mozambique	0.20	0.22	0.34	0.43	0.69
Niger	0.08	0.10	0.18	0.21	0.36
Nigeria	0.09	0.11	0.21	0.22	0.40
Rwanda	0.19	0.23	0.39	0.47	0.88
Senegal	0.08	0.10	0.16	0.17	0.31
Sierra Leone	0.20	0.22	0.34	0.44	0.68
South Africa	0.09	0.12	0.23	0.21	0.44
Swaziland	0.21	0.28	0.60	0.79	1.03
Tanzania	0.15	0.19	0.31	0.36	0.62
Togo	0.09	0.11	0.19	0.25	0.37
Uganda	0.17	0.21	0.41	0.41	0.70

Table 7: Comparing distances between pairs of alternative indexes
(Root mean squared differences over 62 countries of price of consumption.)

	T0	F1.0	F0.5	F0.1	T1.0	T0.5	T0.1	C1.0	C0.5	C0.1
ICP	0.179	0.154	0.155	0.158	0.153	0.156	0.157	0.171	0.176	0.178
F(N)	0.105	0.101	0.104	0.104	0.102	0.106	0106	0.158	0.164	0.167
T(N)	0.093	0.090	0.093	0.093	0.086	0.089	0.090	0.144	0.150	0.153
C(N)	0.107	0.103	0.105	0.105	0.099	0.102	0.102	0.144	0.149	0.152
F(S)	0.073	0.054	0.057	0.057	0.056	0.060	0.060	0.114	0.120	0.123
T(S)	0.073	0.058	0.061	0.062	0.052	0.056	0.057	0.112	0.119	0.121
C(S)	0.084	0.062	0.064	0.065	0.055	0.057	0.057	0.092	0.098	0.102
T0	0	0.062	0.064	0.064	0.058	0.061	0.062	0.121	0.126	0.127
F1	--	0	0.006	0.011	0.023	0.024	0.026	0.075	0.081	0.084
F0.5		--	0	0.010	0.023	0.022	0.024	0.072	0.077	0.080
F0.1			--	0	0.027	0.026	0.026	0.074	0.079	0.081
T1				--	0	0.006	0.012	0.073	0.079	0.082
T0.5					--	0	0.008	0.069	0.075	0.078
T0.1						--	0	0.069	0.074	0.077
C1							--	0	0.011	0.023
C0.5								--	0	0.019
C0.1									--	0

Notes: ICP stands for the price of consumption expenditures by individual households, i.e. the PPP divided by the exchange rate. F(p), T(p), and C(p) are the aggregate (plutocratic) indexes computed from the surveys, Fisher, Törnqvist, and CPD(W) respectively, again divided by the foreign exchange rate. The other indexes are indicated by their first letter, and by the bandwidths in terms of standard deviations of log PCE, 1.0, 0.5, or 0.1.

Table 8: Comparing distances between P4s under different poverty lines
(Means, s.d.'s, and root mean squared differences over 62 countries of price of consumption.)

	Mean	Standard Deviation	Distance from P4 with PL x 2	Distance from P4 with CR PL
Fisher				
Original	1.404	0.379	0.057	0.014
PL times 2	1.455	0.384	0	0.050
CR PL	1.410	0.376	--	0
Törnqvist				
Original	1.402	0.372	0.053	0.013
PL times 2	1.448	0.378	0	0.048
CR PL	1.406	0.372	--	0
CPD				
Original	1.347	0.373	0.101	0.036
PL times 2	1.437	0.381	0	0.084
CR PL	1.362	0.364	--	0

Notes: Original indexes are the prices of consumption based on the P4 index with bandwidth of 0.1 standard deviation; the global poverty line is calculated by weighting by the number of poor people in each of the 50 countries. The PL times 2 uses the same 50 country poverty lines as in the original calculation, but multiplied by two; again, the global line is weighted by the number of people below the line in each countries. This alternative is intended to mimic the comparison between dollar-a-day poverty and two dollar-a-day poverty. The consumption price indexes with CR PL, are intended to mimic Chen and Ravallion's (2008) global poverty line. They are calculated using the poverty lines for 14 of their 15 countries–we do not have data for Guinea-Bissau which is excluded–and without weighting, so that the global poverty line is the unweighted average of the P4 converted value of the 14 lines.

Table 9: Income and regional effects in poverty PPPs versus PPPs and the ICP consumption PPP

	Fisher		Törnqvist		CPD(W)	
	Log of ratio of P4 with bandwidth 0.1 to P3 with NAS weights					
ln *y*	−0.0166	(1.6)	−0.0140	(1.7)	−0.0241	(1.8)
Asia	−0.0202	(0.3)	−0.0206	(0.4)	−0.0366	(0.5)
Africa	−0.0556	(1.0)	−0.0459	(1.0)	−0.0943	(1.2)
Latin America	−0.0275	(0.4)	−0.0218	(0.4)	−0.0351	(0.4)
Western Asia	−0.0353	(0.6)	−0.0429	(0.9)	−0.0826	(1.0)
Constant	0.1280	(1.3)	0.1079	(1.4)	0.1858	(1.1)
F regions (*p*)	0.99	0.42	0.89	0.48	1.75	0.15

Note: India is the omitted "region". The last row shows the *F*-statistic for the omission of the regions, together with the associated *p*-value.

Table 10

Number of poor people in 2005 by region using different poverty lines and purchasing power parity exchange rates (millions)

Global poverty line (international $ or Rupees)		$38	576.86 R	557.00 R	547.83 R	495.06 R	487.94	484.96
# of national poverty lines used		15	14	14	14	50	50	50
PPP type		P3	P4	P4	P4	P4	P4	P4
Aggregation formula		ICP	CPD	Fisher	Törnqvist	CPD	Fisher	Törnqvist
	Population				Number of poor			
World	5,202	1,319	1,209	1,164	1,129	867	874	865
East Asia & Pacific	1,811	308	243	234	231	149	155	159
South Asia	1 451	585	550	516	493	380	370	361
Latin Am. & Caribbean	535	44	42	40	38	31	31	30
East & Central Europe	465	17	14	11	12	9	9	9
Sub-Saharan Africa	698	355	353	356	349	294	306	303
M. East & N. Africa	242	9	6	5	5	3	3	3

Notes: The global poverty line is in terms of monthly per capita expenditure in international dollars or international rupees. The poverty lines for all the P4 versions are calculated simultaneously with the P4s and the poverty counts. The PPP for individual consumption by households from the 2005 ICP for India in international dollars is 15.60.