Poverty PPPs from Latin America and Asia

Angus Deaton, Princeton University
Olivier Dupriez, World Bank

INCOMPLETE DRAFT
The calculations reported here are preliminary, and the results may well change as they are checked and worked on.

March 2007

We gratefully anticipate comments from Yonas Biru, Yuri Dikhanov, Erwin Diewert, Alan Heston, Prasada Rao, and Martin Ravallion.

1. Introduction

This paper presents some preliminary results on the calculation of poverty-weighted purchasing power exchange rates for two groups of countries, four in Asia, and six in Latin America. The prices we use come from the current round of the ICP, while the weights, instead of coming from the National Income and Product Accounts (NIPA) of each country, come from recent household surveys. The availability of household survey data allows us to do two things: (1) recalculate the conventional aggregate purchasing power parity exchange rates using aggregate weights from the surveys, instead of from the national accounts, and (2) to recalculate purchasing power parity exchange rates using weights derived from the expenditure patterns of people at or near the $1-a-day international poverty line. An ultimate aim of the exercise is to generate poverty-weighted purchasing power parity exchange rates that can be used in the calculation of global $1-a-day poverty estimates. These could potentially replace the current PPP rates used for this purpose, whose weights come from aggregate consumption. Because aggregate consumption is the monetary sum of the consumption of all households, richer households get weighted in the total according to their total consumption, so that the aggregate weights may be very different—for example with much less food expenditure—than would be weights derived from the expenditure patterns of households near the poverty line.

We emphasize that the prices we use in our calculations are those from the main ICP, and are identical to those that go into the standard PPPs for consumption in the national accounts. It is sometimes argued that the poor pay different (typically higher) prices than the rich for the same goods, either because they cannot buy in bulk, or because they do not have access to cheaper shops or markets. Whether or not this is the case—and the evidence is mixed to say the least—we make no attempt here to allow for different prices by different economic groups, only different weights. However, there is an important related issue which is that, because the ICP defines consumption to be "final consumption" not "household final consumption," consumption includes government purchases on behalf of households. As a result, the prices that are measured are prices for this concept of consumption, not the prices paid by households. For example, if government buys pharmaceuticals by central purchasing as part of a national health service, the prices recorded by the ICP are the prices paid by the government procurement agency to the pharmaceutical companies, not the prices paid by patients, which are usually subsidized. From the perspective of calculating poverty-related PPPs, these ICP prices and the associated expenditure weights are incorrect. Indeed, commodity subsidies are an explicit poverty-reduction policy in many countries, and our procedures here will not adequately recognize differences in the cost-of-living across countries that are induced by the degree to which governments choose to subsidize prices.

The paper is laid out as follows. Section 2 outlines the procedures but contains only an outline that draws on the main results in Deaton (2006). Section 3 discusses the survey data, and our procedures for matching the surveys—which are heterogeneous and were originally collected for quite different purposes—to the uniform list of prices from the ICP. Section 4 presents results from two groups of countries, Bangladesh, Indonesia, Sri Lanka and Thailand in Asia, and Argentina, Bolivia, Brazil, Columbia, Peru, and Paraguay in Latin America. Section 5 [will one day] conclude and outline tasks for the future.

2. Calculation of poverty-weighted purchasing power parity exchange rates

We shall present three different sets of index numbers, calculated according to the weighted country-product dummy (CPD) method, as well as two EKS-type set of indexes, one based on Fisher ideal index numbers, and one based on Törnqvist price indexes. These calculations follow standard formulas, except that the weights are not the shares of total aggregate expenditures on each good in total aggregate expenditures on all goods, but are instead the shares of each good in the budget of an average person living in a household at or near the poverty line. This calculation is complicated by the fact that the poverty lines in each country must be derived simultaneously with the PPPP indexes. Each country must have the same poverty line in international PPP units—a requirement for calculating global poverty according to a uniform standard—so that we need to know the PPP exchange rates in order to calculate the lines, and we need to know the lines in order to calculate weights for the indexes. We deal with this simultaneity in two ways, first through a closed-form solution that is obtained by assumptions about functional forms, and second, through an iterative procedure that shuttles back and forth between weights and PPPs.

Suppose that we have a set of trial poverty lines, and from them have calculated a set of expenditure weights. For the EKS procedures, we start with matrices of Fisher and Törnqvist indexes, whose element i, j th element is the price index of country j with country i as base. These matrices are then reduced to vectors of PPPs by averaging across countries according to what is known as the EKS method but was first suggested by Gini. The weighted CPD index is computed from a weighted least-squares regression in which the logarithms of the prices of good *i* in country *j* is regressed on a set of country and commodity dummies, with the regression weighted by the expenditure shares. The calculated PPP rates are given by exponentiating the estimated country coefficients. If the base country is omitted from the regression, these are the PPPs with the omitted country as base. Given weights, these procedures yield a set of PPPs. Given the new PPPs, and starting from a poverty line for the base country, we can calculate a new set of poverty lines corresponding to the PPPs, and use them to calculate new weights.

Because there will typically be no households whose per capita household expenditure is *exactly* equal to the poverty line, we need to compute the average expenditure patterns for households *near* the poverty line. This is done by taking a weighted average of the expenditure patterns of households around the poverty line, where the weights are designed so as to give more weight to households the nearer they are to the poverty line. There is a trade-off here between, on the one hand, taking only those households that are close to the poverty line, which focuses on the households we want, but at the risk of not having very many households, and thus imprecise estimates, or on the other hand, taking a wide swathe around the poverty line, which gives greater precision but at the cost of potential bias because we include households that are a long way above or below the line. This trade-off is handled by varying a “bandwidth” which controls the weighting around the poverty line; the smaller the bandwidth, the closer to the poverty line is the focus. Note that, in the current context, below the line is as undesirable as above the line; the exercise here is to calculate PPP rates for converting poverty lines, not for making international comparisons of the living standards of all poor people.

One particular group of formulas is worth rehearsing here because of its relevance to the results that follow, and to the difference between poverty-weighted and non-poverty weighted PPPs. If we focus on two countries, 1 and 2, the Törnqvist index is written

$$\ln P_T^{12} = \frac{1}{2}\sum_{n=1}^{N}(s_n^1 + s_n^2)\ln\frac{p_n^2}{p_n^1} \tag{1}$$

where there are *N* goods, labeled by *n*, p_n^c is the price of good *n* in country *c*, and s_n^c is the expenditure share (or budget share) of the good defined by

$$s_n^c = \frac{p_n^c q_n^c}{\sum_k p_k^c q_k^c} \tag{2}$$

where q_n^c is the quantity of good *n* in country *c*. Given that we are treating the prices as the same for all households, the dependence of (1)—or of any other index—on the level of living depends on how the expenditure shares s_n^c vary across income groups. One simple way to capture this effect is to assume that the Engel curves for the expenditure shares have a specific functional form, and one particularly convenient one is to assume that the budget shares are linear in the logarithm of total household expenditure per capita. Hence

$$s_{nh}^c = \xi_{0n}^c + \xi_{1n}^c \ln x_h + \varepsilon_{nh}^c \tag{3}$$

where the parameters ξ_n^c are country and commodity specific and can be estimated by linear regression within each survey.

Suppose that the poverty line in country 1 is *z*, so that if we are using the Törnqvist index, the poverty line in country 2 is zP_{12}^T so that, by using (3) evaluated at the poverty line in each country, and substituting into (1), we can solve for the poverty-weighted Törnqvist according to

$$\ln P_T^{12} = \frac{0.5\sum_{n=1}^{N}\left(\xi_{0n}^1 + \xi_{0n}^2 + (\xi_{1n}^1 + \xi_{1n}^2)\ln z\ \right)\ln \dfrac{p_n^2}{p_n^1}}{1 - 0.5\sum_{n=1}^{N}\xi_{1n}^2 \ln \dfrac{p_n^2}{p_n^1}} \tag{4}$$

We shall use (4) in its own right, as well as to give starting values for the iterative procedure. As shown in Deaton (2006, Section 5), the procedure can also used where there are more than two countries, with the logarithms of the Törnqvist PPPs given as the solution to a set of linear equations.

However, (4) is of interest in its own right. If the budget shares do not vary with total household expenditure, the parameters ξ_{1n}^c would all be zero, and the term involving z in the numerator of (4), as well as the second term in the denominator, would be zero. In this case (4) is simply the original Törnqvist index, because the ξ_{0n}^c parameters are simply the averages of the budget shares. (Indeed, in this case, they are also the "plutocratic" average of the budget shares, so that the poverty-weighted Törnqvist, the democratically weighted Törnqvist, and the plutocratic or NIPA weighted Törnqvist will coincide.) More broadly, the democratic and poverty-weighted Törnqvist indexes will coincide when parameters ξ_{1n}^c are orthogonal across commodities to the logarithms of the relative prices. Put the other way round, poverty PPPs will differ from democratic PPPs when there is a systematic relationship between Engel effects and relative prices. For example, if we are comparing a rich country and a poor country, and if food in both is mostly traded, then food will be relatively expensive in the poor country. In both countries the food share falls with income according to Engel's Law. In consequence, the move to a poverty-weighted PPP will raise the PPP of the poor country relative to the rich country, raising the local currency equivalent of the international poverty line denominated in the currency of the rich country, and raising the headcount poverty ratio of the poor country. It is harder to think of any such systematic effects between countries at much the same level of development.

All of this holds for the comparison of the poverty-weighted Törnqvist and the democratic Törnqvist whose weights are the averages of the budget shares irrespective of poverty status. However, we are more interested in the comparison between the poverty-weighted PPP and the "conventional" or plutocratic PPP which uses *aggregate* expenditures as weights. In this case, the weights for the index are averages of individual household budget shares weighted by each household's total expenditure. In this case, it turns out that the difference between the poverty-weighted and plutocratic Törnqvist can be written in the form

$$\ln P_T^{12} - \ln \tilde{P}_T^{12} = \frac{0.5\sum_{n=1}^{N}[(\xi_{1n}^1 + \xi_{1n}^2)\ln z - (\xi_{1n}^1 \ln \tilde{x}^1 + \xi_{1n}^2 \ln \tilde{x}^2)]\ln \frac{p_n^2}{p_n^1}}{1 - 0.5\sum_{n=1}^{N}\xi_{1n}^2 \ln \frac{p_n^2}{p_n^1}} \tag{5}$$

where $\tilde{x}^c$ is an (entropy) inequality adjusted measure of mean expenditure

$$\ln \tilde{x}^c = \sum_h \frac{x_h^c}{\sum_h x_h^c} \ln x_h^c \tag{6}$$

and where $\tilde{x}^c$ is measured in common (plutocratic) international prices. These equations tell us that once again, 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. If the countries have the same level of inequality-adjusted income in international units, this condition can be weakened to orthogonality between the logarithms of the relative prices and the two-country average of the ξ_{1n}^c's. 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.

Note that the above argument is specific to the Törqnvist 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 PPPPs to differ from PPPs. The extension to multiple countries is harder to see formally, but our 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 effects, and not by the final adjustment.

We provide two different sets of standard errors for our PPP estimates. The first, and probably the more important, are the sampling standard errors. These treat the prices from the ICP as fixed data, and are concerned with the variability in the indexes that is induced by the fact that the household surveys have only a finite number of observations so that the weights for the indexes are estimated, not known. This is a particular concern for the smaller surveys, and for poverty-weighted indexes in relatively rich countries, where there are relatively few households near the poverty line. The sampling standard errors typically rise when we use smaller bandwidths, so that we are more exclusively focusing on households near the poverty line. They are therefore useful in deciding which bandwidths to use.

We also provide a second type of type of standard error, referred 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

$$\ln p_n^c = \alpha^c + \beta_n + u_n^c \qquad (7)$$

so that, as in the CPD formulation, 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 (7) 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 u_n^c are not zero, and different index number formulae will give different answers for the PPPs. It is this variability across indexes that is captured by the "failure of arbitrage" standard errors. More specifically, the conceptual experiment is one in which the u_n^c are drawn repeatedly, which generates stochastic prices according to (7), which are then combined with non-stochastic expenditure weights to generate stochastic PPPs whose standard errors are calculated. Note that these standard errors are conditional on the expenditure patterns which are taken 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 u_n^c, not directly, but through the PPP indexes.

Once again, the formulae are fully developed in Deaton (2006). Given that the CPD indexes are estimated by running a generalized least squares regression on (7), an estimate of the variance covariance matrix of the estimated parameters can be obtained from

$$\widehat{V}(\widehat{b}) = (X'SX)^{-1}(X'S\Sigma SX)(X'SX)^{-1} \tag{8}$$

where X is the matrix of country and product dummies, S is a diagonal matrix of the expenditure pattern weights, and Σ is an estimate of the variance-covariance matrix of the u_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.

3. Matching surveys to the ICP

Not every country in the ICP has a household survey, and some who do will not allow it to be used for poverty-related analysis. Nor were the available surveys 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.

When the survey categories are finer than those 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 heads, 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 lots of other and potentially more sophisticated ways of synchronizing the two lists, and at some point, it might be worth experimenting with some of them. However, the example of butter and margarine was deliberately 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 available national household survey. In the worst case (here Argentina 1996), weights calculated from the survey will be nine years earlier than the ICP prices (2005) IS THIS RIGHT?. All of the other surveys used here are more recent, with 2002 a modal year. While it would be ideal to be able to match expenditure weights to the year of survey prices, we would expect the expenditure patterns 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.

Another issue occurs is cases where consumption items that are basic heads in the ICP do not appear in the survey. Here we need 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, an imputed item that recognizes the flow of services from houses to their owners who happen also to be their occupiers. 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 can be 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 have made no attempt to add this item back into our surveys. One major concern with any attempt to do so is that rental markets are often primarily 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. Given that we make no correction, the weight on home rental from the surveys (which includes only actual rents) will usually be substantially smaller than the weight in the national accounts; this is possibly an error in the right direction, given that we are trying to construct expenditure weights around the poverty line.

A more extreme case than owner-occupier rents 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 $1-a-day poverty line receive much benefit from these services.

There are also a number of items that are (almost) never represented in the surveys, including purchases of narcotics and prostitution, as well as “purchases by non-residential households in the economic territory of the country.” Together with FISIM, we drop these items from the lists.

There are also items which are included in the ICP but are not purchased in some countries. The most notable examples are pork and alcohol in Muslim countries so that, for example, in our four country Asian calculations reported in the next section, there are no expenditures on these items in Bangladesh, though they appear in Indonesia, Sri Lanka, and Thailand. 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 (7) so that, for example, we impute a price for port in Bangladesh using the country-effect for Bangladesh (which essentially gives us the exchange rate for Bangladesh) and the “pork effects” from the other three countries, which give us a typical relative price for pork. We then leave the item in the survey expenditure files, but assign a zero expenditure to all households.

One final category of expenditures, which is possibly the most important of all, are those that are present in both the national accounts and the surveys, but where the definition is different because the national accounts definition of consumption includes government consumption on behalf of households. This is something about which we can do nothing; it cannot be subtracted out of the national accounts, and we are unwilling to allocate it to households in the survey. Important items here are anything to do with government provided or government subsidized health and education, but the category also includes items such as transportation, where the costs of running the services by the government are larger than the expenditures retrieved from households. Even supposing we could allocate those expenditures and subsidies to individual households—and there exist a number of "benefit incidence" studies of health and education that would help us do so—we are again concerned that such procedures would grossly overvalue these goods, particularly to the poor. Recent work has shown just how common is absenteeism among health-workers and teachers around the world, and has shown just how low service quality can be, particularly for the poor. To attribute health-workers' and teachers' salaries to poor households, and thereby to claim that they are not as poor as would appear, is to add statistical insult to material injury. For these categories of goods and services, we simply live with the fact that the expenditures in the national accounts are very much larger than those in the household surveys, but are not defensive about it, regarding the household estimates as as good or better than those from the national accounts. Indeed, in moving from an aggregate to a poverty-weighted estimate of PPP indexes, the exclusion of the value of these government expenditures is likely to be an important *positive* step.

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 enormous differences. For example, Triplett (199x) 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, Kulshesthra and Kar (2006). Note that we are not here concerned with the increasing divergence in many countries between total expenditures in the surveys and the national accounts. 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 even more 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.

4. Results: PPPs and PPPPs for Latin America and Asia

In this section we report the results of calculating standard and poverty-weighted purchasing power parity indexes for two groups of countries, six in Latin America, Argentina, Bolivia, Brazil, Columbia, Peru, and Paraguay, and four in Asia, Bangladesh, Indonesia, Sri Lanka, and Thailand. The lists of basic heads (consumption categories) are not the same for the two groups, and the relative prices are collected only within the groups, so that it is not possible to pool the ten countries into a single group. All of the results presented here are for the two groups separately. As with the conventional ICP, the poverty PPPs will ultimately be linked across groups using “bridge” countries that appear in more than one group.

The survey data were collected between 1996 (Argentina) and 2003 (Peru and Colombia); of the ten countries only the Argentinian survey was collected before 2000. The sample sizes are given in the first column, second panel of Table 4 for Latin America and in the first column, second panel of Table 7 for Asia; they range from 2,682 in Paraguay to 64,422 in Brazil. Clearly this variation in sample size will affect the sampling errors of the indexes.

We start in Table 1 with the six Latin American countries. The first three columns presents the “conventional” PPP indexes using both prices and expenditures from the ICP. These estimates make no use of the survey data, and are plutocratic indexes because they are based on aggregate expenditures. Columns 1 through 3 present the weighted CPD, EKS-Fisher, and EKS-Törqvist indexes for all ten Latin American countries that are participating in the current round of the ICP. Columns 4 through 6 repeat the exercise for the subset of six countries for which we have survey data, but again using only the ICP expenditures. These two sets of estimates differ only in the final aggregation (or reconciliation) step in the EKS-method (where the matrix of price indexes is turned into a vector), and in the inclusion of additional observations in the weighted CPD regression. The table shows that estimates based on the subset of countries are very close to the estimates from all countries, reflecting something that seems usually to be the case in these calculations, that the uncorrected Fisher and Törnqvist bilateral price indexes are close to satisfying the circularity or path independence property that the price index for country *j* relative to *i* is the same whether we make a one step or many step calculation via one or more other countries. For our purposes here, the importance of the result is that we can eliminate the restriction to a subset of countries as a major source for differences between poverty-weighted and conventionally weighted PPPs. We can hope that this result generalizes to other groups of countries because, even when all of the data are in, we will not have surveys for all of the ICP countries.

The final two columns of Table 1 present the average exchange rates and consumption PPP rate for 2003 from version 6.2 of the Penn World Table, normalized to Bolivia, which is out base country here. (Bolivia is chosen as because it has the highest poverty rate in the region, so that its survey is guaranteed to have a substantial number of households at or around the international poverty line.) Apart from Peru, all of the PWT rates are lower than those from the current round of the ICP, suggesting that it is the Bolivian (and Peruvian) rates that are the major difference between the PWT and the current ICP.

Table 2 presents the same plutocratic PPPs, but now with expenditure weights derived by aggregating up the surveys, and without use of the expenditures from the national accounts. As will consistently be the case, the two EKS indexes are very close to one another, but neither differs greatly from the weighted CPD index. If we focus (arbitrarily) on the Fisher indexes, and compare Tables 1 and 2, the Peruvian and Brazilian PPPs (relative to Bolivia) are essentially identical. The Argentine PPP is about six percent lower from the surveys, the Brazilian PPP about 3.5 percent lower, the Columbian PPP about 3 percent lower, and the Paraguayan estimate is 8.2 percent lower, which is the largest discrepancy. [In the next version of the paper, we will work to trace back these discrepancies to the differences in weights that generate them. In the meantime,] Table 3 shows, for each of the six countries, the five basic heads whose expenditure shares differ by the largest amounts in absolute value. For each country, the table also shows the correlation between the shares derived from the surveys and from the national accounts; note that we are comparing like with like here, because the survey shares were calculated from estimated aggregate expenditures on each good. These correlations range from 0.86 for Brazil to 0.36 for Peru although, apart from Peru, all are larger than 0.58. The goods with the largest discrepancies are typically those that we would expect. Actual and imputed rentals of dwellings is the top category in five of the six countries though, surprisingly, it is not true that the national accounts figure is always larger than the survey figure. It appears that in Brazil, Colombia, Peru and Paraguay, but not in Argentina or Bolivia, imputed rents are added into the survey, though it is not clear how this was done nor why, if this is the case, the two estimates would not be closer. Education, health (including pharmaceutical products), and passenger transportation appear as expected, with larger shares in the national accounts than in the surveys (except for pharmaceuticals in Bolivia.) There are other categories, such as garments, where there is no obvious reason for a discrepancy. [**We clearly need to do more work here and understand better these discrepancies and perhaps even make further adjustments to the surveys to ensure a more uniform treatment. These discrepancies may be as or more important than poverty weighting in moving from PPPs to PPPPs.]**

Table 2 also shows estimated standard errors (of both types) for the PPPs calculated from the surveys. The first point to note is that the sampling standard errors are very small; the largest (for the weighted CPD for Brazil) is less than half of one percent of the estimate. The two EKS indexes appear to make more efficient use of the survey data, and their sampling standard errors are substantially less than those of the weighted CPD index, sometimes only half as large. As we shall see below, the virtual negligibility of the sampling standard errors is something that carries through to the poverty-weighted PPPs, even with small bandwidths, and even in the countries with the smallest surveys. Because the survey expenditure shares are calculated over all survey households (or all households near the poverty line) including those who do not purchase the good, and because the PPP indexes are further averaged over the hundred or so basic headings of consumption, sampling variability is very small.

The same is not true of the "failure of arbitrage" standard errors, also reported in Table 2. These are typically between five and ten percent of the estimated PPPs; of course, this is not a measure of estimation uncertainty in the usual sense, but a measure of the extent to which relative prices differ across countries, and the consequences for the ambiguity of an exchange rate measure.

Table 4 presents the poverty-weighted PPP indexes for the six Latin American countries. We have used Bolivia as the base country and chosen a poverty line of 1,710 bolivianos of annual per capita expenditure in 2002; this generates a headcount ratio of 24 percent which matches the World Bank's estimate for that year. At the end of the calculations, each of the other five countries will have a poverty line that is the calculated PPPP multiplied by the Bolivian poverty line and those poverty lines, in turn, are generating the expenditure weights for the indexes.

The top panel of Table 4 starts from the Törnqvist approximation that can be directly calculated, and then shows the iterated Törnqvist indexes for a range of bandwidths expressed in terms of the standard deviation of the logarithm of per capital total household expenditure. The first such estimate uses a bandwidth of a full standard deviation, the second 0.5 of a standard deviation, then 0.1, and finally 0.05. As the bandwidth decreases, the number of households in the band around the poverty line also decreases so that, taking Argentina as an example, where there are 27,275 households in the survey, we get 11,611 households included at a bandwidth of one standard deviation, 5,542 at a half, 1,083 at a tenth, and 530 at a twentieth. The second panel shows these numbers for all six surveys; for Paraguay, which has the smallest sample size, there are only 59 households in the band around the poverty line when the bandwidth is 0.01 standard deviations. The final panel of the table shows the sampling standard errors which increase as the bandwidth decreases; even so, and even when there are only 59 observations, the relative sampling standard error is never larger than one percent, almost certainly negligible given the other errors and uncertainties in the estimation of these indexes.

Table 3 does not report the "failure of arbitrage" standard errors. These are very similar to those listed in Table 2, and do not vary much either across index types are across the different bandwidths.

The last two columns of the Table show the Fisher EKS and weighted CPD indexes for a bandwidth of 0.1; this is not the smallest bandwidth but we found that, in the case of Paraguay, the weighted CPD index looked implausible with the smallest bandwidth, and this is not surprising given the larger sampling variance of the CPD index (see also Table 2.) We do not report the full range of bandwidths for these indexes because, like the Törnqvist indexes, they change very little with the bandwidth. Instead, we have selected 0.1 standard deviations as the smallest bandwidth (and presumably least bias) for which there are no obvious signs of unacceptable sampling imprecision. The largest relative standard error is now 1.24 percent for the weighted CPD index.

As to the estimates of the PPPPs themselves, they are remarkably close to the conventional (survey-based) PPPs in Table 2. While it is not entirely clear what is the appropriate metric for looking at differences—and calculating headcount ratios can sometimes turn small differences in PPPs into large difference in poverty rates—the largest of the differences (using the EKS-Fisher with a bandwidth if 0.1 standard deviations) is for Colombia, where the Bolivia-based PPPP is 4.9 percent smaller than the PPP in Table 2. For Argentina, the PPPP is 2.7 percent larger than the PPP, for Brazil it is 2.3 percent less, for Peru 2.8 percent less, and for Paraguay, 1.3 percent larger. While these numbers are certainly larger than the sampling standard errors (about twice as large), they are much smaller than the "failure of arbitrage" standard errors which are a better measure of our uncertainty about international PPPs. And it seems likely that statistical errors and discrepancies of measurement are likely more important in practice than the conceptual issue of whether or not the indexes are poverty weighted.

The reason that the poverty-weighted PPPPs in Table 4 and the plutocratic PPPs in Table 2 are so close is that there is, in practice, little correlation between the effects of total expenditure on the budget shares, the ξ_1 parameters, on the one hand, and the relative prices, on the other, see equations 4 and 5 above. Again, if we look at this on a pairwise basis with Bolivia as base, and inspect the terms in equation (4), the denominators are all close to unity (1.04, 0.99, 1.02, 0.98, and 0.98) while the numerators are dominated by the first term (which involves the intercepts of the budget share equations, the ξ_0's) rather than the second term (which involves the effects of log expenditure on changing the budget shares, the ξ_1's. For Argentina, the two terms in the (log) Törnqvist approximation are –0.842 and 0.115, for Brazil –0.553 and –0.036, for Colombia 6.05 and 0.06, for Peru –0.47 and –0.00, and for Paraguay 6.48 and 0.12. In these data, there is not enough of a relationship between the structure of relative prices and the way in which expenditure shares vary across income groups to drive a large wedge between the poverty-weighted and conventional PPPs.

The final tables of the paper are for the four Asian countries for which we currently have data, Bangladesh, Indonesia, Sri Lanka, and Thailand. Here, we take Bangladesh as base, and also choose the Bangladeshi headcount ratio to be held fixed in the calculations. A poverty line of 6,712 taka per capita per annum gives a headcount ratio in 2000 of 36.04 percent, which is the official World Bank estimate. In this case, we do not have the parities and the national accounts based expenditure shares for countries other than the four, so we start with Table 5, which shows the conventional PPPs together with the PPPs for 2003 from the Penn World Table 6.2, and Table 6, which shows their survey-based equivalents. Taking the EKS-Fisher indexes to compare, the survey based estimates are 4.3 percent, 4.8 percent, and 2.4 percent larger than the national accounts based estimates for Indonesia, Sri Lanka, and Thailand respectively. Clearly, much of the difference could be reassigned to Bangladesh, which is the base. As was the case for Latin America, the sampling standard errors appear to be negligibly small although, once again, the weighted CPD index has much larger sampling uncertainty than either of the two EKS indexes. For the latter, the largest sampling standard error is only 0.13 percent, a tiny amount relative to other uncertainties. The failure of arbitrage standard errors are a good deal smaller than for Latin America, presumably reflecting less international variation in relative prices, or perhaps simply the smaller number of countries.

Table 7 presents the poverty-weighted PPPs for the four countries with Bangladesh as base. The table follows the same format as Table 4 for Latin America, and the qualitative results are very similar. Sampling standard errors remain small, and there is relatively little variation in the results across index type or bandwidth. Most importantly, the PPPPs in Table 7 are extremely close to the survey-based PPPs in Table 6. Again using the EKS-Fisher to illustrate, the Indonesian PPPP is 1.3 percent less than the PPP, the Sri Lankan PPPP is 0.75 percent less than its PPP, and the Thai PPPP is 0.97 percent higher than its PPP. While these differences are large relative to the sampling standard errors, it is only because the standard errors themselves are so tiny. As was the case with Latin America, the similarity of the PPPP rates to the PPP rates can be traced back to lack of any correlation between income variation in expenditure patterns, on the one hand, and the relative prices on the other.

Table 8 reports the poverty headcount ratios calculated using the baseline poverty lines for Bolivia and Bangladesh respectively, converted to other poverty lines using the calculated PPPs, in all cases using the EKS-Fisher to illustrate. The three columns are calculated from (1) the conventional PPPs using the aggregate expenditure weights from the national accounts, (2) the conventional PPPs using the aggregate expenditure weights estimated from the surveys, and (3) the poverty-weighted PPP rates. As would be expected from the previous discussion of the PPPs, the differences along the rows are not very large, and tend to be larger between the first and second columns than between the second and third. As before, the measurement issues, and the definitions of consumption, appear to be more important than the poverty-weighting.

Table 1

PPP exchange rates for LAC: six countries versus ten countries: aggregate weights from NIPA

Countries	PPP: NIPA aggregates			PPP: NIPA aggregates: six countries only			PWT6.2	
	CPD	Fisher	Törnqvist	CPD	Fisher	Törnqvist	XRAT	PPP_c
Argentina	0.535	0.519	0.519	0.540	0.525	0.526	0.379	0.419
Bolivia	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Brazil	0.660	0.641	0.638	0.612	0.599	0.595	0.402	0.494
Colombia	461.2	456.4	454.6	463.4	457.0	457.0	375.8	408.5
Peru	0.663	0.652	0.649	0.622	0.615	0.612	0.454	0.798
Paraguay	835.7	818.2	809.2	833.3	822.9	811.8	838.8	611.1
Chile	155.4	152.0	151.7	--	--	--	90.27	124.4
Ecuador	0.202	0.198	0.198	--	--	--	0.131	0.219
Uruguay	6.196	6.100	6.101	--	--	--	3.747	5.100
Venezuela	498.7	488.0	486.6	--	--	--	209.8	434.0

Notes: The PWT6.2 indexes are the product of “pc” and “xrat” and are for 2003, normalized to Bolivia. XRAT is the average exchange rate from the same source.

Table 2

PPP exchange rates for LAC: six countries: aggregate weights from surveys

		CPD			Fisher			Törnqvist		
	Year	PPP	s.e.(1)	s.e.(2)	PPP	s.e.(1)	s.e.(2)	PPP	s.e.(1)	s.e.(2)
Argentina	1996	0.484	0.39	5.41	0.494	0.19	6.01	0.494	0.19	6.01
Bolivia	2002	1.000	--	--	1.000	--	--	1.000	--	--
Brazil	2002	0.609	0.43	10.8	0.590	0.21	9.08	0.591	0.21	9.08
Colombia	2003	441.4	0.31	5.62	443.2	0.20	4.72	443.3	0.20	4.72
Peru	2003	0.605	0.26	6.27	0.604	0.21	5.32	0.605	0.20	5.32
Paraguay	2000	758.7	0.35	6.83	765.3	0.23	6.48	767.3	0.21	6.48

Notes: Year in the second column is the year of the household survey used to construct the weights. Standard errors are shown as *relative* standard errors, or standard errors of the logarithm of the estimate of the PPP, expressed in percentages so that, for example, the actual standard error for Brazil is 0.003 (0.609 × 0.0043). The standard errors are labeled (1) for sampling standard errors, and (2) for “failure of arbitrage” standard errors, see text for explanation.

Table 3
Expenditure shares (percentages) from surveys and national accounts

Category	Survey share	NAS share
Argentina $\rho=0.686$		
Rentals actual & imputed	4.80	15.46
Education	2.36	5.28
Hospital Services	0.19	2.48
Maintenance of dwellings	2.45	0.23
Cultural service	2.48	0.29
Bolivia $\rho=0.587$		
Passenger transport by road	2.65	10.47
Education	0.80	8.46
Fuels for personal transport	1.06	4.65
Pharmaceuticals	4.70	1.15
Garments	5.52	2.28
Brazil $\rho=0.861$		
Rentals actual and imputed	20.55	13.77
Education	3.43	9.59
Motor cars	6.11	3.33
Maintenance of dwelling	2.08	0.11
Pharmaceuticals	2.85	4.62
Colombia $\rho=0.754$		
Rentals actual and imputed	17.69	10.12
Education	2.30	8.39
Food products n.e.c.	5.39	0.69
Insurance	4.14	1.10
Pharmaceuticals	0.84	3.69
Peru $\rho=0.361$		
Rentals actual and imputed	17.28	4.70
Education	1.25	9.31
Catering services	0.17	7.91
Passenger transport by road	0.18	5.41
Garments	0.43	4.56
Paraguay $\rho=0.678$		
Rentals actual and imputed	14.30	5.79
Garments	0.62	5.58
Education	0.41	5.05
Catering services	7.06	4.01
Motor cars	0.19	2.85

Table 4

PPPP exchanges rates: Poverty weighted: Törnqvist approximation, and fixed point estimates at various bandwidths

	Törnqvist Indexes					Fisher	CPDW
Bandwidth	Approx.	1	0.5	0.1	0.05	0.1	0.1
Argentina	0.461	0.484	0.483	0.485	0.486	0.481	0.465
Bolivia	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Brazil	0.518	0.605	0.603	0.603	0.604	0.604	0.639
Columbia	461.7	465.9	455.7	468.1	467.7	466.2	471.0
Peru	0.612	0.618	0.618	0.620	0.620	0.621	0.620
Paraguay	751.7	761.4	761.4	759.1	759.1	755.1	752.8
Numbers of observations within bandwidth around poverty line (first column is total number of households in the survey)							
Argentina	27,275	11,611	5,542	1,083	530	1,083	1,024
Bolivia	5,732	3,178	1,642	323	168	323	323
Brazil	48,564	26,096	13,655	2,863	1,421	2,840	2,961
Columbia	22,979	7,097	3,148	551	257	554	555
Peru	18,911	10.788	5,572	1,089	576	1,094	1,088
Paraguay	2,682	1,258	659	128	59	127	127
Estimate of relative standard error from sampling, percentages							
Argentina		0.153	0.191	0.359	0.497	0.438	0.92
Bolivia		--	--	--	--	--	--
Brazil		0.217	0.271	0.520	0.749	0.569	1.01
Columbia		0.141	0.173	0.310	0.406	0.317	1.00
Peru		0.128	0.152	0.271	0.360	0.298	0.95
Paraguay		0.179	0.236	0.591	0.998	0.108	1.24

Table 5

PPP exchange rates for Asia: four countries: aggregate NIPA weights

Countries	PPP: NIPA aggregates			PWT6.2	
	CPD	Fisher	Törnqvist	XRAT	PPP_c
Bangladesh	1.000	1.000	1.000	1.000	1.000
Indonesia	148.4	148.9	149.5	147.5	201.3
Sri Lanka	1.576	1.529	1.556	1.750	1.920
Thailand	0.700	0.701	0.700	0.713	1.281

Table 6

PPP exchange rates for Asia: four countries: aggregate weights from surveys

		CPD			Fisher			Törnqvist		
	year	PPP	s.e.(1)	s.e.(2)	PPP	s.e.(1)	s.e.(2)	PPP	s.e.(1)	s.e.(2)
Bangladesh	2000	1.000	--	--	1.000	--	--	1.000	--	--
Indonesia	2002	150.8	0.14	7.38	155.2	0.12	6.38	153.8	0.10	6.38
Sri Lanka	2002	1.560	0.20	7.11	1.603	0.12	6.12	1.604	0.10	6.12
Thailand	2002	0.707	0.21	6.35	0.718	0.13	5.79	0.712	0.12	5.79

Notes: Year in the second column is the year of the household survey used to construct the weights. Standard errors are shown as *relative* standard errors, or standard errors of the logarithm of the estimate of the PPP, expressed in percentages. The standard errors are labeled (1) for sampling standard errors, and (2) for "failure of arbitrage" standard errors, see text for explanation.

Table 7

PPPP exchanges rates for four Asian countries: Poverty weighted: Törnqvist approximation, and fixed-point estimates at various bandwidths

	Törnqvist Indexes					Fisher	CPDW
Bandwidth	Approx.	1	0.5	0.1	0.05	0.1	0.1
Bangladesh	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Indonesia	156.1	152.8	152.7	152.5	152.4	153.2	146.3
Sri Lanka	1.617	1.593	1.591	1.591	1.589	1.591	1.536
Thailand	0.695	0.720	0.721	0.721	0.720	0.725	0.696
Numbers of observations within bandwidth around poverty line (first column is total number of households in the survey)							
Bangladesh	7,448	5,149	2,996	614	325	614	614
Indonesia	64,422	23,064	10,692	2,011	995	2,048	1,742
Sri Lanka	16,924	5,847	2,564	488	245	489	442
Thailand	34,785	2,720	625	77	46	85	72
Estimate of relative standard error from sampling, percentages							
Bangladesh	--	--	--	--	--	--	--
Indonesia	--	0.11	0.15	0.30	0.35	0.33	0.38
Sri Lanka	--	0.11	0.14	0.24	0.32	0.26	0.34
Thailand	--	0.14	0.22	0.48	0.57	0.47	0.61

Table 8

Headcount ratios at alternate PPP rates (percentages)

	PPP National Accounts	PPP Aggregate survey weights	PPP Poverty-weighted
Argentina	18.55	16.73	16.12
Bolivia	24.00	24.00	24.00
Brazil	18.13	17.69	18.34
Colombia	14.26	13.62	14.81
Peru	20.41	19.65	20.80
Paraguay	28.95	26.89	26.25
Bangladesh	36.04	36.04	36.04
Indonesia	5.22	6.50	6.09
Sri Lanka	3.50	4.58	4.39
Thailand	0.25	0.29	0.31

Notes: All calculations based on EKS-Fisher indexes.