

NBER WORKING PAPER SERIES

TRYING TO UNDERSTAND THE PPPS IN ICP2011:
WHY ARE THE RESULTS SO DIFFERENT?

Angus Deaton
Bettina Aten

Working Paper 20244
<http://www.nber.org/papers/w20244>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2014

We are grateful to Orley Ashenfelter, Grant Cameron, Erwin Diewert, Bill Easterly, Rob Feenstra, Nada Hamadeh, Alan Heston, Robert Inklaar, Michel Mouyelo-Katoula, Paul McCarthy, Prasada Rao, Marko Rissanen, Marcel Timmer, and Fred Vogel for help, comments and suggestions. The views expressed in this paper are solely those of the authors and not necessarily those of the U.S. Bureau of Economic Analysis, the U.S. Department of Commerce, or the National Bureau of Economic Research. Both Aten and Deaton are consultants to the International Comparison Program (ICP); Aten as a member of the Computation Task Force, and Deaton as a member of the Technical Advisory Group. The research reported in this paper was not funded by the ICP. The data used in the paper are either from the World Development Indicators or can be accessed via application to the ICP; for details see

http://siteresources.worldbank.org/ICPINT/Resources/270056-1255977254560/121120_ICPDataAccessPrinciples&Procedures.pdf

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Angus Deaton and Bettina Aten. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Trying to Understand the PPPs in ICP2011: Why are the Results so Different?

Angus Deaton and Bettina Aten
NBER Working Paper No. 20244
June 2014, Revised April 2015
JEL No. E01,F00,O11,O47

ABSTRACT

Purchasing power parity exchange rates, or PPPs, are price indexes that summarize prices in each country relative to a numeraire country, typically the United States. These numbers are used to compare living standards across countries, by academics in studies of economic growth, particularly through the Penn World Table, and in some cases, to allocate resources. The International Comparison Program (ICP) collects the detailed prices on which these indexes are based on an irregular basis. In 2014, the ICP published PPPs from the 2011 round that are sharply different from those that were expected from extrapolation of the previous round, ICP 2005. These discrepancies will eventually have important implications for the Penn World Table, and for international comparisons of living standards. The world according to ICP 2011 looks markedly more equal than calculated from ICP 2005. This paper investigates why this happened. We identify a likely source of the problem in the way that the regions of the ICP were linked in 2005. We use two different methods for measuring the size of the effect. Both suggest that the 2005 PPPs for consumption for countries in Asia (excluding Japan), Western Asia, and Africa were overstated relative to the US by between 18 to 26 percent.

Angus Deaton
361 Wallace Hall
Woodrow Wilson School
Princeton University
Princeton, NJ 08544-1013
and NBER
deaton@princeton.edu

Bettina Aten
Bureau of Economic Analysis
1441 L Street, N.W.
Washington, DC 20230
Bettina.Aten@bea.gov

1. Introduction

Purchasing power parity exchange rates, PPPs, are international multilateral price indexes that measure, for the various components of GDP, the amount of local currency required to purchase the same real amount in that country relative to a numeraire, which is typically the United States. Non-traded goods are typically cheaper in poorer economies, so that PPPs are typically lower than exchange rates for poor countries, and are more so the poorer the country: for example, in 2011 the market exchange rate for India was 46.7 rupees to the dollar, while the PPP exchange rate for consumption was 15.0 rupees to the dollar. PPP-adjusted national accounts appear in several widely used data sets, such as the Penn World Table and the World Development Indicators. The European Union uses PPP-adjusted accounts for redistributive purposes, while adjusted GDP numbers enter into the IMF quota formulas. The World Bank uses PPP exchange rates to calculate global poverty estimates. The burgeoning interest in global income inequality also requires such numbers.

The PPP indexes are based on thousands of prices collected in nearly 200 countries. The work is done by the International Comparison Program (ICP), which has collected prices in a series of benchmark years, the latest of which are 2005 and 2011. A summary of the results of the 2011 round was published in April 2014, World Bank (2014); the report tabulates purchasing power parity exchange rates for the main aggregates of GDP for 199 economies. Until the publication of these results, the World Bank in its World Development Indicators provided extrapolated PPP exchange rates; for most countries, these were based on the 2005 round of the ICP, updated using relative inflation rates for each country. The new estimates for

2011 from ICP 2011 are quite different from these extrapolations. In particular, most poor countries of the world are estimated to be larger relative to the US and other rich countries than was estimated from the extrapolations. This aspect of the results has attracted a good deal of attention, particularly the fact that the aggregate Chinese economy is much closer to the US than previously estimated, and also that the Indian economy is now estimated to be larger than the Japanese economy. The new results also sharply reduce previous estimates of international inequality.

The PPPs for individual consumption by households, like the PPPs for GDP, have been revised downwards, so that estimated consumption levels outside the rich countries are now higher than previously estimated. As was immediately noted by several commentators, Dykstra, Kenney, and Sandefur (2014), and Chandy and Kharas (2014), these consumption PPPs are the relevant ones for the calculation of global poverty rates, which are likely to be deeply affected. If the global poverty line is held fixed in real 2005 US dollars, for example at \$1.25 as currently used by the World Bank, its local value in poor country currencies will now be lower, so that there will be fewer people living below it. These calculations show very large declines in estimated poverty rates for 2010—which is the last year covered by the World Bank’s poverty calculator. Chandy and Kharas calculate that the number of people living below \$1.25 in 2010 has fallen from 1,215 million using the PPPs extrapolated from ICP 2005, to 571 million using the new PPPs back cast from 2011 to 2010. Poverty in India falls by more than 300 million people. These are estimates by outside commentators not the World Bank, and are conditional on the line being held at \$1.25 at 2005 prices.

The new revisions will ultimately work their way into the Penn World Table, PWT, which will not only be updated, but will likely see substantial revision of past years. Much more modest revisions of the PWT have, in the past, changed the conclusions of academic studies based on them, Johnson, Larson, Papageorgiou and Subramanian (2013), so that it is important that researchers understand the origins of the current revisions. This paper aims to provide a non-specialist account, while presenting an analysis in support of our conclusions.

The 2011 round of the ICP contained many methodological improvements over ICP 2005 but even in the absence of methodological changes, the results from a new round will not be identical to those extrapolated from an earlier round. Extrapolation is at best a short cut to the collection of new benchmark data. Even so, it remains unclear exactly why the 2011 results are quite so different from the extrapolations and, without such an understanding, there will be continuing questions about whether the new results really are better than the old, and whether it is safe to use them.

The 2011 ICP was not the first time that a new round of the ICP brought large changes to PPPs and to poverty counts. ICP 2005 saw large *increases* in poor-country PPPs relative to the US compared with those that had been previously extrapolated from the 1993/95 round. These changes caused a sharp *upward* revision of the number of poor people below any global line that is denominated and held fixed in real US dollars. There were also upward revisions in 1993/95 compared with the previous round in 1985; these also had large effects on the poverty calculations, Deaton (2001). When we think about what happened in 2011, we must keep

this previous history in mind. In particular, if ICP 2011 is inconsistent with ICP 2005, the problems could come from either ICP 2005 or ICP 2011, or both.

In this paper, we focus on the PPPs for individual household consumption that are used to compare living standards between different countries. This narrowing of focus also simplifies the discussion by taking off the table some (although not all) of the most difficult (“comparison resistant”) items, such as construction or government services, as well as the trade-balance, which is treated differently in ICP from national accounts, and so predictably causes discrepancies between ICP estimates of GDP and extrapolations based on national accounts, McCarthy (2011). A more ambitious treatment that recalculates ICP 2005 using ICP 2011 methodology, and treats all of GDP is Inklaar and Prasada Rao (2015).

In section 2, we document the changes in the consumption PPPs in 2011 and look back on the corresponding changes in 2005. These calculations demonstrate that, to an extent, the ICP 2011 undid some of the changes that happened in the ICP 2005. In section 3, we discuss three possible explanations for the conflict between extrapolations and actuals, two general, and one specific. The first general point is that extrapolation cannot be expected to match actual changes, essentially because the ICP is based on international multilateral price indexes that are different from the national price indexes that are used for extrapolation. The second general point is that PPPs have large, but largely unrecognized, standard errors, so that at least some of the changes with each round likely reflect random measurement error. Finally, and more specifically, we check whether there is evidence that there is a problem with the national prices indexes used for extrapolation.

In section 4, we document our ultimately preferred explanation, which is the change in the way that the regions were linked in 2011. The ICP is organized on a regional basis, with each region calculating a set of within-region PPPs. In order to get a world table, these regions have to be linked. In 2005, this was done using a “ring” of eighteen countries, distributed over the regions, for whom there was a separate ICP exercise. The PPPs for those countries from the ring comparison were then used to calculate four regional price indexes that scaled the regional PPPs to give global PPPs. At the time of the 2005 ICP, there were concerns about the implementation of the ring method—for example, some of the within-region ring comparisons were inconsistent with the same comparisons from the main ICP, and some items on the ring list were difficult to price in some of the ring countries. In consequence, the ring method was not used in 2011. Instead, a global core list was developed, and *all* countries priced (country-specific subsets of) the items on the list together with local items that were specific to the region. The overlaps in the global list allowed all the calculations to come from a single set of data, and *all* countries in all regions contributed to the regional linking in a symmetric way.

Deaton (2010) suggested that the ring in ICP 2005 may have contributed to over-pricing in poor countries relative to rich countries, so that it is possible that ICP 2011, with its global core list, undid this overpricing. We provide some evidence from the ring countries that is consistent with this hypothesis. If this is correct, the main problem lies with ICP 2005, not ICP 2011. Our estimates suggest that countries in the African, Asian, and Western Asian region of the ICP had consumption PPPs that were overestimated by 18 to 26 percent in 2005.

Section 4 also provides an alternative approach based on an examination of the relationship between price levels and per capita consumption. According to the Balassa-Samuelson theorem, price levels—the ratios of PPPs to exchange rates—are higher in richer countries and rise as countries become richer. Ravallion (2012) suggested that the relationship between price levels and per capita consumption is stable enough to use for extrapolation and gives a better fit than CPI extrapolation, a finding challenged by Inklaar (2013). Here, we use the relationship for a different purpose, to assess whether price levels in 2005 were systematically overstated. This method is far from ideal if only because the Balassa-Samuelson effect is unlikely to operate uniformly across just two rounds, and because fluctuations in exchange rates move PLIs around in a way that can obscure the relationship except in the long run. Even so, our estimates from this method are consistent with those from examination of the ring, and suggest an overestimation of around 18 percent in Africa, Asia, and Western Asia in ICP 2005.

Our results provide a plausible story for the conflict between the 2005 and 2011 ICP benchmarks (as well as for the conflict between the 1993/96 and 2005 benchmarks), but the evidence is far from conclusive. There are many other changes from ICP 2005 to ICP 2011 that the ring hypothesis does not explain; given the many other things that changed, that is not in itself surprising. Even so, we note that the results for consumption in Inklaar and Prasada (2015), who use yet another method to check for ring price overstatement, come to estimates that are close to those in this paper.

2. Explanations

To understand what happened in ICP 2011 we must go back to ICP 2005, and how those earlier data differed from the extrapolations based on the previous (1993/95) round. Figure 1 plots the ratio of the “new” (i.e. ICP 2005 PPP) to “old” (i.e. extrapolated from 1993/95) PPPs for individual household consumption in 2005 against the logarithm of per capita GDP in 2005 international prices. All data relate to 2005 and the GDP figures are those taken from the report of the 2005 round of the ICP. We include countries (such as China and India) that were not part of ICP 1993/95 and whose PPPs were imputed in that round. Given that the imputations in 1993/95 were in line with benchmarked countries and typically estimated from the relationship between PPPs and per capita GDP for the benchmarks, Figure 1 looks similar if we only use countries present in both benchmarks

The immediate feature of Figure 1 is that the 2005 revisions were mostly upward, i.e. the ratios are greater than one. The countries on the right that have a ratio of close to one are the Eurostat-OECD countries, who publish annual numbers, and therefore see only minimal revisions at new rounds of the ICP; we return to the Eurostat-OECD case below. Among the other countries, there is a pronounced negative slope, so that the lower was per capita GDP, the larger the upward revision in the consumption PPP. Relative to the US, poor countries got poorer or, equivalently, relative to most of the world, the US and other rich countries got richer; for example, if India has been the numeraire, and PPPs calculated in local currency per rupee, the line would have sloped upward to the right. Either way, the 2005 ICPs made the world distribution of income more unequal compared with extrapolations from the

previous round. Some of the revisions are very large, mostly in sub-Saharan Africa, but there were also upward revisions in Asia, including factors of 2.6 for Philippines, 2.3 for China, 2.1 for Indonesia, 1.9 for Nepal, and 1.4 for India and for Bangladesh.

When the 2005 results were published, these revisions were noted, but at the time, there seemed no pressing need to explain them. ICP 2005 had introduced many methodological improvements compared with previous rounds, some of which were expected to cause large revisions, especially given the widespread concern about the quality of ICP 1993/95, see Ryten (1998). Many new countries had been surveyed for the first time in 2005, including many in sub-Saharan Africa where the largest revisions had taken place. China officially participated in the ICP for the first time in 2005 and India did not participate between 1978 and 2005. It was therefore reasonable to conclude that the new results were simply better, and that the revisions were largely a consequence of the poor quality of ICP 1993/95.

Figure 2 replicates Figure 1, but compares ICP 2011 with extrapolations based on ICP 2005; again the Eurostat-OECD numbers come from the annual program and are not extrapolated. It shows the ratio of consumption PPPs from ICP 2011 to the extrapolations for 2011, plotted against per capita GDP in 2011 international prices. Figure 2 is to a limited extent the opposite of Figure 1. The ratios of new to old are now mostly less than one instead of greater than one. Just as countries got poorer relative to previous calculations in the 2005 round, they got richer in the 2011 round. Even so, the pronounced negative slope in Figure 1 does not turn into a positive slope in Figure 2. But the graph is somewhat misleading because it includes, on the bottom right, a number of wealthy non-OECD countries, most of

which are oil producers and whose per capita GDP levels are thereby inflated relative to consumption. If those countries are excluded, the positive slope is apparent. As we shall see in section 3 below, perhaps the most important methodological improvement in 2011 over 2005 is capable of explaining these results, not country by country, but for regions as a whole.

If we compare the log ratios of extrapolated to actual in 2005 and 2011, the correlation is -0.39 , which is entirely driven by the comparison between the Eurostat-OECD countries and the rest of the world. ICP 2011 reversed some of the increase in inequality between poor and rich regions in ICP 2005, but it did not reverse the 2005 changes on a country-by-country basis. Recall that ICP 1993/95 was a weak round, with many imputations, so that many of the 2005 ratios come from problems with that earlier round and, as we shall see, the variance of the ratios was much lower in 2011 than in 2005.

Table 1 shows the revisions for each of the ICP regions, again covering only those countries that were covered in both rounds. The revisions in 2005 relative to extrapolations were very large, particularly in Africa, though all of the rest of the world had higher PPPs relative to the rich countries. These changes made the world distribution of consumption appear to be radically more unequal. Depending on the region, about 40 percent of those increases were undone in 2011, though in Western Asia, the downward revision in 2011 was larger than the upward revision in 2005. The reduction in Africa in 2011 undid much less of the 2005 increase than was the case in India, so that, just as the ICP 2005 “Asianized” poverty, the ICP 2011 will “Africanize” it, at least if no offsetting changes are made to the global poverty

line; this effect is amplified by the fact that there are many more people close to the (or any plausible) global poverty line in South Asia than is the case in Africa.

Table 1 also shows the standard deviations of the log ratios of extrapolations to actuals. Apart from the CIS—among which there are many special problems—there is a marked reduction in within regional dispersion in 2011 compared with 2005. This is what we would expect from the fact that the methodology and country coverage was much more stable between 2005 and 2011 than before 2005. A much larger share of the changes between 2005 and 2011 comes from between regional revisions than from within regional revisions. This does not tell us which of 2005 and 2011 is correct, only that, within the regions, they have more in common than with previous rounds. The finding is also consistent with our hypothesis that the main reason for the differences between 2005 and 2011 is to do with the way that the regions were linked.

3. Why are new ICPs not consistent with extrapolations?

3.1 Spatial and temporal price inconsistency in general

It has long been known that spatial and temporal price indexes cannot be consistent with one another even under ideal conditions, with perfect measurement. There are also many practical considerations that can widen the gap, see McCarthy (2012) for a recent summary and cites to the literature. Indeed, if extrapolations were consistently accurate, there would be no need for the ICP at all. Extrapolations, although intuitively attractive, cannot be expected to reproduce the numbers that come out of an ICP benchmarking exercise.

The PPPs from the International Comparison Program are *multilateral* price indexes, in which the PPP for each country within a region depends on price relatives and weights from all of the countries in the system. The difference between multilateral and bilateral price indexes is present even when there are only two countries, and this simplest of cases can be used to illustrate the issues involved.

When the World Bank estimates a PPP by extrapolation, it starts from a baseline consumption PPP, for example from ICP 2005, and then updates using consumer price indexes. If the PPP for India in USD were 20, say, and if CPI inflation in both countries were the same, the extrapolated PPP would remain at 20. If Indian CPI inflation were higher than CPI inflation in the US, the extrapolated figure would be higher than 20, and if the Indian inflation were lower than inflation in the US, it would be lower than 20. Of course, domestic CPIs use only local prices and local weights, whereas the PPP, when it is calculated by the ICP, will combine weights for both India and the US in calculating a PPP between them. Multilateral price indexes start from bilateral indexes between pairs of countries, and if those bilateral indexes are to satisfy minimal properties, for example, that the Indian to US PPP must be the reciprocal of the US to Indian PPP, then they must use weights from both countries. As a result, the change in the PPPs by extrapolation will be different from the changes from one ICP to the next.

In the two-country case, with countries 1 and 2, one simple formula for the relative change in CPI in 2 relative to 1 is

$$d \ln P_2 - d \ln P_1 = s_2' d \ln p_2 - s_1' d \ln p_1 \quad (1)$$

where P_2 and P_1 are the two country price indexes, p_2 and p_1 are the underlying vectors of prices of goods and services, and s_2 and s_1 are the two corresponding vectors of weights or budget shares. The PPP for 2 in terms of 1, by contrast, uses weights for both countries, so if we use a Törnqvist index, we have

$$\ln PPP_2 = 0.5(s_2 + s_1)'(\ln p_2 - \ln p_1) \quad (2)$$

If we combine these two to see what happens over time, we have, ignoring changes in weights,

$$d \ln PPP_2 = (d \ln P_2 - d \ln P_1) - 0.5(s_2 - s_1)'(d \ln p_2 + d \ln p_1) \quad (3)$$

The last term is the discrepancy between the extrapolation, which is the first term on the right, and the change in the PPP on the left hand side. Its magnitude will depend on the changes in the underlying relative prices, as well as differences in the consumption patterns in the two countries. The discrepancy is zero if there are no relative price changes in either country, if the structures of spending are the same, or if the differences in spending structures are uncorrelated with average changes in relative prices. For countries with very different patterns of consumption, or with different structures of domestic inflation, the discrepancy could be large. For example, if the relative price of food rises around the world—as was the case from 2005 to 2011—the change in the benchmark PPPs for a poor country relative to a rich country is predicted by equation (3) to be less than the extrapolated value. This line of enquiry is worth pursuing further, but calculations by the ICP team suggest that the effect was not large relative to the extrapolation discrepancies discussed here; calculations by Inklaar and Timmer (2013) for the Eurostat-OECD countries show

that, even in countries with strong statistical offices, the discrepancies are much larger than can be explained by equation (3).

These results assume that the prices measured in the CPI are identical to the prices measured in the ICP, which is far from being the case. Cross-country indexes must match goods that are (a) identical in both locations, (b) reasonably commonly consumed in both places, so that the comparison lists for the ICP are usually different from the comparison lists of the CPI. Once again, it is hard to assess what this would contribute to the discrepancies.

3.2 Quality of local CPIs, and the increase in food prices

Another concern is the quality of domestic CPI indexes, which varies considerably across countries; common problems are outdated weights, political interference in a very visible indicator, or a coverage that may be restricted to urban consumers, or even only to those who live in the capital city. CPIs are also relevant when thinking about why price level indexes (PLIs) differ between poor and rich countries, and how PLIs can be expected to change over time. CPIs cover *all* prices, including both traded and untraded goods, as opposed to exchange rates which are determined by flows of traded goods; indeed, the fact that exchange rates do not reflect the prices of untraded goods is the principal reason for collecting PPP data in the first place. As countries get richer, and wages rise, the price of untraded goods, which are initially much lower than in poor than in rich countries, will rise and, unless obscured by changes in the exchange rate, the PLI will rise too, as all prices converge to world prices, see Obstfeld and Rogoff (1996, pp. 210–4) for a textbook account. Because the CPI covers both traded and untraded goods, and provided it is correctly meas-

ured, it will automatically rise relative to the exchange rate as the country grows richer. Ravallion (2012) calls this the dynamic Penn effect in a reference to the original Penn effect, by which poorer countries have lower prices relative to their exchange rates. If the CPI is correctly measured, extrapolation forward from a previous PPP should automatically capture any such effect, at least up to the other reservations already discussed in this section. In the absence of problems with CPIs, the existence of a dynamic Penn effect does not invalidate the forward extrapolation of PPPs using CPIs.

Between 2005 and 2011, there was a large worldwide increase in the relative price of food in world markets; the FAO food price index doubled in nominal terms from 2005 to 2011, and rose by 50 percent in relative terms. In countries whose CPI weights are outdated, the share of food in the index will be too high—provided the country is growing and the food share falling—and the measured CPI growth will be too high. Because the extrapolation is done using CPIs, and if, as is plausible, CPI weights are more likely to be current in rich countries, the extrapolation would lead to consumption PPPs that are too high for poor countries relative to rich countries, which would be revealed at the time of a new PPP.

One way of assessing this possibility is to compare inflation rates from CPIs with inflation rates from the implicit price deflators of consumers' expenditure in the national accounts. This is not perfect, because the coverage of the CPI and consumption in the national accounts will not be the same, but the implicit price deflator—current price divided by constant price consumption—is a current weighted Paasche index, while CPIs are generally Laspeyres indexes, so the effect of rising rel-

ative prices of food overstating the latter is likely to be indicated by a divergence in the two rates of inflation.

There are some countries where, indeed, CPI inflation between 2005 and 2011 was considerably higher than inflation in the price deflator. India is one example, where CPI inflation was 50.4 percent and deflator inflation was only 38.3 percent. But this is not generally the case. Chinese CPI inflation was 19.6 percent, nearly 11 percent *lower* than the 30.4 percent inflation in the deflator. The World Development Indicators contain 115 countries with both indexes and, of these, only 61 have CPI inflation greater than deflator inflation. There are two countries, Nicaragua and Mauritania, where CPI inflation exceeds deflator inflation by more than 0.50, and three countries, Venezuela, Tajikistan, and Tanzania where the deflator exceeds CPI inflation by more than 0.50. Venezuela is a country where the CPI may have been politically controlled, while the much less visible implicit price deflator was not. The remaining countries are shown in Figure 3, which plots the difference against the logarithm of per capita GDP in (the base year) 2005 in 2005 international dollars.

The Figure does not suggest that the food price story is the reason that extrapolation gave PPPs that were consistently too high, if only because the points are not consistently above the zero line for the poorer countries. Yet the wide divergence between the two rates of inflation, outside of the rich countries, reinforces concerns about the quality of price data in many countries. And even if these problems cannot explain the fact that PPPs were consistently lower than extrapolated, they surely contribute to the variance that is observed in Figure 2 and Table 1.

3.3 Standard errors of PPPs

One issue that is not always given sufficient recognition is that PPPs are themselves subject to substantial uncertainty in a world where relative prices and consumption patterns differ substantially across countries. Deaton (2012) shows that the logarithm of the ratio of the Laspeyres to the Paasche index is a measure of uncertainty about the PPP between any pair of countries, and derives formulas for uncertainty in multilateral indexes. The idea underlying these calculations is to write the price of good i in country c using the country-product decomposition formula

$$\ln p_{ic} = \alpha_i + \beta_c + \varepsilon_{ic} \quad (4)$$

where α_i is a commodity fixed effect, β_c is a country fixed effect, and ε_{ic} is the residual from the projection of log price on the commodity and country factors. The factor β_c can be thought of as the logarithm of a (simplistic) PPP exchange rate, while α_i is the typical relative price of good i ; a liter of cooking oil is generally more expensive than a liter of milk, or a kilo of rice, and so on. Equation (4) would be exact in a world in which the law of one price holds true.

Standard errors are calculated by thinking of equation (4) as a data-generating mechanism, and calculating how random draws of ε affect the ultimate PPP formulas. The resulting relative standard errors are large, 20 to 30 percent for US to India or China comparisons. These are derived for a single international exercise, and the errors in changes may be smaller if the sources of error are correlated over time; indeed it is plausible that ε_{ic} 's, although highly variable across countries,

because relative prices vary from country to country, are relatively stable over time, if goods that are relatively expensive or cheap in a country are persistently so. Food subsidies in India would be one example.

There is much work to be done in this area; but what we know now suggests that differences in relative prices across countries, and over time within countries, can generate substantial random variation in calculated PPPs.

3.4 Discrepancies in the EU-OECD region

It is useful to consider the last row in Table 1, which shows the discrepancies that would have occurred in the Eurostat/OECD countries if, contrary to actual practice which updates PPPs annually, the 2011 PPPs had been predicted using CPIs, as was the case for the rest of the world. The average discrepancy (growth of PPP minus growth of CPI) is -4.3 percent over six years, or -0.72 percent a year and is much smaller than in any other region; the mean absolute value of the deviations is still a modest 5.7 percent. As noted above, Inklaar and Timmer (2012) find that equation (3) cannot account for the discrepancies, and it remains unclear what drives them. The discrepancies are not correlated with per capita GDP, but are higher in places where the rate of inflation is higher, a result that is even stronger when we look at the discrepancies between PPP growth and the implicit price deflator of GDP, which is used to update PPPs for GDP between and beyond rounds. High rates of inflation tend to come with changes in relative prices, which are the fuel for discrepancies in indexes. Beyond that, of the ten largest discrepancies, eight are in formerly communist countries, the others being Turkey and (more surprisingly) Britain.

These Eurostat-OECD numbers are comparable in magnitude to the discrepancies in the United States across cities (technically metropolitan statistical areas) between inflation measured from multilateral spatial price indexes and inflation in local CPIs, Aten and Reinsdorf (2010).

These mostly rich country comparisons provide a useful calibration for what we might expect for the rest of the world. Discrepancies in indexes are more likely when inflation rates are high, and where patterns of consumption and relative prices are markedly different. Yet, with high quality statistical offices and no changes in method, the results suggest that the discrepancies are typically small. The Eurostat-OECD region is a (favorable) outlier in Table 1, which suggests that, for the other regions, discrepancies come from some methodological change between ICP2005 and ICP2011.

4. Regional linking, the ring and the comparison between 2005 and 2011

4.1 Regional linking methods in 2005 and 2011

One methodological difference between the 2005 and 2011 rounds of the ICP was the way in which the regions were linked to give a full global set of international PPPs. The ICP is organized into regions, each of which runs its own set of price comparisons to calculate PPPs for each of the countries within its region. Each of these regional ICPs has its own numeraire so, at this first stage, it is not possible to compare, for example, the United States, which is in the Eurostat/OECD region, with Cameroon, which is in the African region. At a final stage, a set of price indexes is calculated, one for each region, and these are used to scale the regional parities, re-

specting the original relative PPPs within regions, and to reach a full international comparison. It is this last step that was done differently in 2011 from 2005. Because linking moves entire regions, it is an obvious place to look if we suspect that all of Asia, or all of Africa, was too low or too high in one of the rounds.

In 2005, regional linking was done using eighteen “ring” countries, selected so that there were at least two in each region; these are listed in the first column of Table 2. The ring methodology was not seen as ideal, and was developed as a response to a number of constraints in the 2005 round, some of which were not apparent at the planning stage; there was no prior empirical experience in using the method. By the time of the linking, which comes at the end of the process, most of the regions had completed their own intraregional PPPs, and those were respected in the global analysis. For Eurostat/OECD, this “fixity” of within-region PPPs is a precondition for participation in the broader ICP, and in 2005 was imposed for all regions. This procedure has the advantage of preventing countries with very weak data, or that are far away from another region, from affecting the calculation of the PPPs within any region to which it does not belong. The selection of ring countries was also limited by the willingness of countries to undertake a second, different, pricing exercise.

The ring worked through a separate price comparison exercise for the ring countries alone, ignoring the rest of the world, and using a specially constructed list of more than 1,000 goods and services. These PPPs for the eighteen ring countries were then used to derive regional price indexes that could be used for linking. The details of the method, and the formulas, are given in Diewert (2013) or in Deaton

and Heston (2010). It starts by selecting a numeraire country in each region, and then converting all the regional prices into the currency of that region's numeraire. In this way, each region becomes a "super-country" to be linked using the regional price indexes from the ring. After the round was complete, it was discovered that the super-country method is not invariant to the choice of numeraire country within each region; calculations by the ICP team suggest this is not a major source of error, but it was one factor in the decision to abandon the ring in 2011.

When the 2005 ICP was completed and reviewed, a number of other problems with the ring method became apparent. Some of the ring countries were very different from one another, for example Cameroon and Japan, or Brazil and Oman, making comparisons difficult, and raising the concern that the selection of ring countries might be having a large effect on the PPPs of all countries. Britain is the only representative of Western Europe (or, apart from Japan, of the OECD), while there are four representatives in sub-Saharan Africa. Beyond that, some of the within-region comparisons from the ring turned out to be inconsistent with the comparisons from the previously constructed within-regional comparisons for the main calculations. The latter, which used regional lists adapted to each region's pattern of consumption, and which covered all countries in the region, were almost certainly superior to the ring comparisons, which were based on a common list, with some items hard to price, and for only a handful of countries in the region. Yet it was only these ring comparisons that determined the "tectonic" price indexes that linked one region with another.

Given these difficulties, all of which were apparent before the next round was planned, ICP 2011 used a common global "core" list, developed in consultation with all of the regions. Countries in each region priced as many of the items on the list as they could together with additional items selected by the region as suitable for their countries. The idea was not for all countries to price all the items on the core list, which would not be feasible, but to provide enough overlap so that, at the final linking stage, all countries would be involved, there would be many links across regions, and the goods and prices would be the same as those that had been used in the construction of the within-region PPPs at the first stage. The precise linking method in 2011 is the Country Aggregation with Redistribution (CAR) method, first introduced by Kravis, Summers, and Heston (1982) in the 1980 ICP; a brief description is given in the Summary Report, World Bank (2014). The advantage of the CAR method over the ring lies in its consistency with the regional pricing strategy and in its incorporation of *all* countries into the linking of the regions so that we escape the dependence of the all-important regional linking factors on the choice of ring countries, or on the selection of goods and services for the ring list.

4.2 Did the 2005 ring overstate price levels in poor countries?

That the ring method might have overstated PPPs in poor countries was explored in Deaton (2010) though, in the absence of the 2011 ICP, there were no decisive results. The ring list was extremely detailed, with a large number of items that were only widely available in rich countries. Enumerators in Cameroon, Kenya, Senegal, Zambia and Sri Lanka had to price a 2003 or 2004 vintage bottle of Bordeaux, front-

loading washing machines with a pre-specified spin speed, and a Peugeot 407 with air conditioning and climate control. Although all of these goods are traded, they are not commonly purchased in poor countries, they are often subject to high taxes, and would be high priced in any case because quantities are too small to overcome fixed costs of importation or domestic production. The obvious danger is that, when such items can be found at all, they are rare and expensive, or perhaps only available in stores patronized by a small subset of the population; in the absence of expenditure weights for items within basic heads there will be an overstatement of the price level in poor relative to rich countries. (Basic heads are the building blocks of categories of consumption, and are the lowest level at which expenditure weights are available, so that the prices of red Bordeaux, along with other wine items, enter into the basic head of wine.)

A second and related problem with ring pricing is what happens when a consumption category is very expensive but rarely consumed in the poor country, but relatively cheap and widely consumed in the rich country. Air travel is an example that is relatively very expensive in many African countries. It has a low budget share in the African ring countries, so that its high price there has little effect on the overall domestic price level. But when we compare air travel in Cameroon, say, with air travel in Britain, say, the high relative price in Cameroon is weighted, not by the low Cameroon share, but by an average of the shares in Britain and Cameroon. This makes Cameroon look very expensive, even though hardly anyone there buys the good, and given the small number of countries in the ring, this high price of Came-

room relative to Britain can have a substantial effect on the price index for Africa relative to the OECD.

Questions of item specification—which can be thought of as quality issues—and of how to weight—are not specific to the ring method, but are inherent in making multilateral price comparisons across widely different countries. Yet it is plausible that the 2005 ring, with its very detailed list with many rich-country items, and the limited number of countries, led to an overestimation of price levels in the poorer countries in the ring, and so overestimated price levels for whole regions, particularly for Asia and for Africa.

We investigate the “ring overstatement hypothesis” by looking at how the ring countries fared in both 2005 and 2011. We calculate price levels for consumption for the ring countries using ICP data for 2005 and 2011, ignoring all other countries, so that we have comparable ring mini-ICPs for both rounds; the 2005 version is the same as that used to link the regions, while the 2011 version is only for the purposes of this paper. We then use the two comparisons to calculate an ICP-based inflation rate from 2005 to 2011 for each ring country, which we can compare with the inflation rate from the CPI. If the ring led to price overstatement in poor countries in 2005, the ICP-based inflation rates will be systematically lower than CPI inflation in poor countries, at least up to the noise that is apparent in Figure 3.

For 2005, we used the ring prices and ring list to calculate a set of PPPs, as if the ring countries were the only countries in the world. In 2011, we have done the same thing with the global core list, but again only using information from the eighteen countries that comprised the ring in 2005. No information from any other coun-

try has any effect on these numbers. For both rounds, we used the country-product dummy (CPD) method below the basic head level with weights reflecting surveyor-reported importance of each item (for 2011, unweighted in 2005), and expenditure-weighted CPD above it. Ideally, we should have liked to calculate price indexes in local currency for each country and each basic head, so as to make a direct comparison between 2011 and 2005. However, for many basic heads there is little overlap between the ring list in 2005 and the core list in 2011, so this approach is not feasible.

Table 2 lists the results. The PPPs and the exchange rates are scaled so that the UK is the numeraire. The relative inflation rates are changes in the logs of the CPIs from 2005 to 2011 from the World Development Indicators minus the corresponding inflation rate of the UK. We have adjusted the Estonian and Slovenian exchange rates and PPPs for 2011 back to their original currencies by multiplying by the rate at the time they joined the Euro.

If extrapolation were to work perfectly, the change in the logs of the PPPs from 2005 to 2011 would be equal to the inflation rate in the last column. We already know that this is not going to happen. Even so, if for some countries the ring prices were overstated in 2005, but not in 2011, we would expect the change in the PPPs to be systematically *less* than the inflation rate because, in addition to true inflation, the overstatement would be unwinding. We think of the log change in the PPP as the predicted or ICP-based inflation rate, to be compared with the actual inflation rate, and what we are looking for is that the ICP-based inflation rate will be

too low in the poorer countries of the ring, and approximately correct in the richer countries.

The vertical axis of Figure 4 shows the residuals between ICP and CPI inflation for the ring countries, where the ICP inflation comes from the ring exercise in 2005 and our new calculations for the ring countries only in 2011. All numbers are normalized relative to the U.K., whose ICP and CPI inflations are therefore zero, as is its residual. The horizontal axis is the logarithm of per capita GDP in 2005 prices from the full ICP 2005, and the scatter shows that the discrepancy is negative for the poorest countries, rising to zero for Britain and the rich countries, as is to be expected if the ring prices were too high in the poorer countries in the ring in 2005, but were correctly measured in 2011. If we (somewhat arbitrarily) take Malaysia, Senegal, Cameroon, Philippines, Oman (excluded from the graph for lack of an income figure), Jordan, Zambia, Sri Lanka, Egypt, and Kenya as the poor group, for whom the ring overstatement was an issue, the average discrepancy for those ten countries relative to the rich group is 0.35. Note that, although South Africa (ZAF) was in Africa and in the ring, the wide availability of rich country goods there means that it should be where we see it, with the rich countries. Japan has markedly different consumption patterns than most other rich countries, and so it is plausible that there would be some ring overstatement there.

We realize that we are reading a great deal into 18 data points. We also note that there is a good deal of variation around the line, but, subject to country to country variation in ring pricing in 2005, country to country variation in the accuracy of the CPI, and all the other reasons why PPPs are not expected to inflate at the same

rate as the CPI, it is perhaps surprising that the relationship in Figure 4 is as strong as it is. Given the questionable quality of the consumer price indexes in some countries, we may be asking too much for the growth of the PPPs to match one-for-one with the growth in the CPIs. Table 3 runs regressions of the change in log PPP on the change in log CPI, without constraining the coefficient to be one. The first two columns show separate regressions for the two groups; the average distance between the two lines is 0.26. The last column pools the 18 countries, imposing the same slope on both groups, and calculates the distance between the two lines as the coefficient on a dummy for the poor countries, which also gives 0.26. A more relevant number comes from putting South Africa into the poor group, given that it was part of the Africa region in ICP 2005. This reduces the distance between the lines to be 0.22 (two regressions) or 0.23 (one regression, results not shown.) We prefer these more conservative estimates, but have no formal account of why the PPPs should not move one-for-one with the CPIs.

If the analysis here is correct, and given that the ring prices moved whole regions relative to OECD/Eurostat, the \$ PPPs for Asia, Africa, and Western Asia were all too high in 2005, by 22 percent according to our best estimate. This number should be compared with the revisions in the second column of Table 1; if the 2005 ICPs for the three regions were reduced by this amount, the extrapolations to 2011 would have been close, not on a country-by-country basis, but region by region.

Inklaar and Rao (2015) propose an alternative calculation that uses the fact that in the ring countries in 2005, there were two pricing exercises, one for the original regional PPPs using the lists for each region, and one for the ring calculation,

using the ring list. This calculation makes no use of the 2011 data, and is therefore different from the calculation here. It also has the advantage that the overstatement can be calculated for each country, although only relative to a regional numeraire. Comparing with South Africa in Africa and Hong Kong in Asia, they find an average overstatement in Africa of about 9 percent and in Asia of 16 percent. Figure 5 suggests that Hong Kong had no ring price overstatement, while South Africa is 16 percent below Britain and Hong Kong, so the final average overstatement is 25 percent (16+9) for Africa, and 16 percent for Asia. These estimates compare with the estimate above of 22 percent for Africa, Asia and Western Asia together.

4.3 A crosscheck calculation

A different approach is to look for evidence of overstatement in the cross-country relationship between the price level index (PLI)—the ratio of the PPP to the exchange rate—and per capita consumption. The Balassa–Samuelson theorem posits that PLIs are lower in poor countries, and if this relationship is stable over time—which may or may not be the case—and if the PPPs of Africa, Asia and Western Asia were overstated in 2005, we should see an upward shift in the relationship for those regions in that year. Although we report results for this method, we regard it as unsatisfactory. The dynamic Balassa-Samuelson works through time, and the most appropriate way to test it would be to follow at least one country over a substantial period of time. In a cross-section, or in a short time series of cross-sections like those from successive rounds of the ICP, the relationship will be obscured by time series effects, of which the most important are likely to be those associated with

currency movements. Currencies can deviate from fundamentals for long periods of time, and non-fundamental movements of major currencies, such as the dollar or the euro, can distort the cross-sectional relationship between PLIs and per capita GDP in any given year, or over a run of years. That caveat should be kept in mind in reading the rest of this section.

Suppose that the relationship between the *PLI* in country *c* and the level of economic development can be approximated by an equation that is linear in logs

$$\ln PLI_c^* = \alpha + \beta(\ln C_c^* - \ln C_0) + u_c \quad (5)$$

and that this holds both across countries and across time, at least for the correctly measured C^* (in *PPP* terms) and PLI^* . C^0 is per capita consumption in *PPP* for a reference country, such as the United States, or the geometric mean of per capita consumption for a group of countries, such as the Eurostat–OECD countries; it is measured without error. The Balassa–Samuelson theorem is silent about exactly which measure of resources should represent the level of development on the right-hand side of (5). Our choice of a relative measure is motivated by the fact that, in the numeraire country, in our data the US, always has a *PLI* of unity, so that, if the US economy grows between two rounds, it is impossible for it to lie on the same curve if the resource variable is an absolute measure. The relative specification (5) gets round this problem, but there is no guarantee that it is correctly specified. Indeed, since relative prices of goods change from one round to the next, the relative measure in (5) can change with them, even if there has been no change in a country’s real position. We will use 2005 international dollars for the ratios in ICP2005, and 2011 international dollars for the ratios in ICP2011. Note also that (for the same reasons)

the estimates of (5) are *not* invariant with respect to the choice of the comparison country 0; we will consider two possibilities here, the US and a Eurostat-OECD composite. These ambiguities are yet another reason (in addition to exchange rate fluctuations) to be wary of the calculations reported in this subsection.

Suppose that for three regions, Western Asia, Africa, and Asia, and in one year only, 2005, the *PPPs* are overstated, so that

$$\ln PPP_c = \ln PPP_c^* + \theta_R \quad (6)$$

where θ_R is a positive number for the selected regions in 2005, and zero otherwise.

Note that $\ln PLI$ and $\ln C - \ln C_0$ are also derived from the mismeasured *PPP* and are thus subject to a corresponding downward shift of θ . If we substitute (6) into (5), we get the relationship between the measured *PLI* and measured *C*, which is

$$\ln PLI_c = \alpha + \theta(1 + \beta) + \beta(\ln C_c - \ln C_0) + u \quad (7)$$

so that equation (7) should show an intercept shift in 2005 for the three regions.

Figure 5 shows the scatter of the logs of the *PLIs* and log of relative per capita consumption in both ICP 2005 and ICP2011, with dark points for 2011 and lighter points for 2005. India and China are shown as larger circles, and the changes between rounds by connecting arrows. Both ICPs have an upward slope, but it is clear that, on the left of the figure, among the poorer countries, the lighter circles tend to be above the darker circles at the same level of per capita consumption. Note too that there is a good deal of scatter; for example, although China got richer and its *PLI* increased, India, which also got richer, saw its *PLI* decline. Of course, if the ring overstatement hypothesis is correct, both of those changes are biased downwards because their *PLIs* in 2005 were too high.

To assess the degree of overstatement in ICP 2005, we run Balassa-Samuelson regressions (7). We do so for countries in the three regions (Africa, Asia, and Western Asia) only; this enables us to calculate estimates of θ from the intercept shift between the two rounds. It also has the advantage that we are not assuming that the relationship is linear across rich and poor countries alike, which Figure 5 suggests would not be appropriate. The results are shown in Table 4, on the left hand panel for the US as reference country, and on the right hand panel for the geometric mean of Eurostat-OECD as reference “country.” In the first two columns of each panel we estimate separate regressions for the two years, and then average the difference between the two regression lines over all countries. To get an estimate of θ we must divide by the slope of the regression; when we use the average over the two, we get an estimate of *PPP* overstatement of 18 percent in those three regions. The third column shows the simpler case where we allow only an intercept difference, which yields an estimate of θ of 21 percent. When we use the compound country as reference, the corresponding estimates are 21 percent and 19 percent. These estimates are somewhat lower than our preferred estimate of 23 percent estimated in the previous section but well within the margins of error we should expect.

Neither of these calculations is very precise, and both make assumptions about functional forms and the absence of other effects. Even so, all of our evidence is consistent with overestimation in 2005 for Asia, Africa, and Western Asia, by an amount of around 20 percent.

5. Summary and conclusions

We began by documenting the differences between, on the one hand, the PPPs for consumption that were measured in the 2011 round of the ICP and, on the other hand, their extrapolations based on the 2005 round together with relative rates of inflation using country consumer price indexes. There are good reasons to expect the extrapolations to differ from actuals, but we show that, in the countries of the OECD and Europe whose PPPs are calculated by Eurostat, the divergence is small, even using the extrapolation procedures used for the rest of the world. Outside this region, and especially in Africa, Asia, and Western Asia, the divergences are large.

ICP 2005 revised upward the PPPs for poor countries relative to the rich countries and ICP 2011 revised them downward. Country-by-country, the revisions in 2011 are only weakly correlated with the comparable revisions in 2005, and in aggregate, only remove part of them. In 2005, when there were large methodological differences for many countries, the discrepancies with the extrapolations from the previous round in 1993/96 varied both across regions and within them. In 2011, by contrast, where the methodological revisions were more limited, the biggest revisions were across regions, not within them.

The pattern of raising poor region PPPs in 2005, and lowering them in 2011 raises the possibility that there was a one time problem in 2005 that was, at least to some extent, undone in 2011. A plausible explanation for this is that the regions were linked in 2005 using a group of eighteen ring countries, a method that was replaced by a more robust and internally consistent procedure in 2011, a method that also involved all countries, not just a small group. Our direct analysis of the ring

countries, comparing the inflation rates from one ICP to the next, and comparing it with inflation rates from national CPIs, suggested that the PPPs for the Western Asia, Africa, and Asia regions might have been inflated by around a fifth to a quarter. A different approach, based on an analysis of the Balassa-Samuelson or Penn effect in 2005 and 2011, which has problems of its own, is nevertheless consistent with a similar overestimation in 2005. This last method is conditional on assuming that 2011 is correct, and 2005 not, so it is less strong than the method that looks at inflation in the ring countries. But both sets of results are consistent with the up and down pattern shown by the discrepancies between actuals and interpolations in 2005 and 2011.

It is not our purpose in this paper to discuss how these results should be used. But our findings suggest that the ICP 2011 estimates are the most accurate that we have, and provide no grounds for doubting them. Indeed, there are other major improvements in 2011 that we have not discussed, including much more rigorous validation of data, and much more attention to the local currency national accounts supplied by the countries. If this is correct, then the revisions that need to be undertaken are to long-standing previous estimates, a process that is likely to be less than straightforward.

Given our results, one way to correct 2005 would be to preserve the within-region PPPs for each of the regions in 2005, but to revise the regional linking. This could be done by back casting (backward extrapolation from 2011 to 2005) for each country and then aggregating GDP and consumption up to regions. The regional totals would then be allocated to the countries in proportion to the original 2005 es-

timates, which would preserve the within-region totals. This is essentially the Country Aggregation with Redistribution (CAR) method that was used to link the regions in 2011. As far as the future is concerned, it would also be extremely desirable to put the program on a continuous basis, with results updated from year to year, which should avoid the large revisions that characterized the last two rounds.

6. Citations

Aten, Bettina, and Marshall Reinsdorf, 2010, "Comparing the consistency of price parities for regions of the U.S. in an economic approach framework," Presented at the 31st annual conference of IARIW, St Gallen, Switzerland, August.
http://www.bea.gov/research/topics/price_indexes.htm

Chandy, Lawrence, and Homi Kharas, 2014, "What do new price data mean for the goal of ending extreme poverty?" Brookings, May 5,
<http://www.brookings.edu/blogs/up-front/posts/2014/05/05-data-extreme-poverty-chandy-kharas>

Deaton, Angus, 2001, "Counting the poor: problems and possible solutions," *World Bank Research Observer*, 16:2, 125–47.

Deaton, Angus, 2010, "Price indexes, inequality, and the measurement of world poverty," *American Economic Review*, 100:1, 5–34.

Deaton, Angus, 2012, "Calibrating measurement uncertainty in purchasing power parity exchange rates," incomplete paper,
http://siteresources.worldbank.org/ICPINT/Resources/270056-1255977254560/6483625-1338834270350/ADeaton_CalibratinguncertaintyinPPP.pdf

Deaton, Angus, and Alan Heston, 2010, "Understanding PPPs and PPP-based national accounts," *American Economic Journal: Macroeconomics*, 2:4, 1–35.

Diewert, Erwin, 2013, "Methods of aggregation above the basic heading level: linking the regions," Chapter 6 in Fred Vogel and D. S. Prasada Rao, *Measuring the real size of the world economy: the framework, methodology and results of the international comparison program*, Washington, DC. The World Bank, pp. 169–96.

Dykstra, Sarah, Charles Kenny, and Justin Sandefur, 2014, "Global absolute poverty fell by almost half on Tuesday," Center for Global Development, May 2, <http://www.cgdev.org/blog/global-absolute-poverty-fell-almost-half-tuesday>

Inklaar, Robert, 2013, "Price levels and economic growth: making sense of revisions to data on real incomes: a comment," *Review of Income and Wealth*, 59:4, 614–22.

Inklaar, Robert, and D. S. Prasada Rao, 2014, "Cross-country income levels over time: did the developing world suddenly become much richer?" Groningen Growth and Development Center Research Memorandum 151, Groningen, The Netherlands, December.

Johnson, Simon, William Larson, Chris Papageorgiou, and Arvind Subramanian, 2013, "Is newer better? Penn World Table revisions and their impact on growth estimates," *Journal of Monetary Economics*, 60(2), 255–74.

Kravis, Irving, Alan Heston, and Robert Summers, 1982, *World product and income: international comparisons of real gross product*, Baltimore, MD. Johns Hopkins University Press for the World Bank.

McCarthy, Paul, 2011, "Extrapolating PPPs and comparing ICP benchmark results," International Comparison Project, http://siteresources.worldbank.org/ICPINT/Resources/270056-1255977254560/6483625-1291755426408/18_ICPBook_Extrapolating_FINAL.pdf

Obstfeld, Maurice and Kenneth Rogoff, 1996, *Foundations of international macroeconomics*, MIT Press.

Ravallion, Martin, 2012, "Price levels and economic growth: making sense of revisions to data on real incomes," *Review of Income and Wealth*, 59:4, 593–613

Ryten, Jacob, 1998, *Evaluation of the International Comparison Programme*, United Nations Economic and Social Council, New York, United Nations.

World Bank, 2014, *Summary of results and findings of the 2011 international comparisons program*, International Comparison Program, <http://siteresources.worldbank.org/ICPINT/Resources/270056-1183395201801/Summary-of-Results-and-Findings-of-the-2011-International-Comparison-Program.pdf>

Table 1: Regional revisions from extrapolations in ICP 2005 and ICP 2011

Region	Mean log ratio 2005	Mean log ratio 2011	s.d. log ratio 2005	s.d. log ratio 2011	Number of countries
Africa	0.56	-0.25	0.27	0.15	33
Asia & Pacific	0.46	-0.34	0.27	0.08	15
CIS	0.40	-0.15	0.24	0.25	5
Eurostat/OECD	0.08	-0.00	0.11	0.01	40
Latin America	0.26	-0.13	0.21	0.06	7
Western Asia	0.42	-0.56	0.14	0.12	5
Eurostat/OECD extrapolated		-0.04		0.06	45

Notes: CIS is the Confederation of Independent States. The means are calculated only for countries that are common to both rounds, and are unweighted country means of the logarithm of the ratio of the ICP consumption PPP to the previously extrapolated consumption PPP. The final row is what would have happened in 2011 if the Eurostat/OECD PPPs had been projected from 2005 using consumer price indexes; of the 47 Eurostat/OECD countries, Cyprus is excluded, and Chile was not an OECD member in 2005. Extrapolation can be done for more countries than have full PPP data for all three rounds, so that there are 45 countries in Eurostat/OECD in the last row, as opposed to 40 earlier in the table.

Table 2: Data on PPPs, exchange rates, and inflation from 2005 to 2011

Country	PPP 2005	PPP 2011	XR 2005	XR 2011	Relative Inflation rate
Brazil	2.801	2.720	4.418	2.679	0.114
Chile	684.1	641.0	1018.4	780.8	-0.132
Cameroon	649.8	404.2	959.0	756.2	0.003
Egypt	3.570	3.275	10.51	9.474	0.465
Estonia	15.51	13.64	22.87	18.03	0.104
UK	1	1	1	1	0
Hong Kong	9.764	9.128	14.15	12.47	-0.017
Jordan	0.6786	0.533	1.291	1.138	0.158
Japan	218.2	160.37	200.4	127.9	-0.186
Kenya	64.80	61.77	137.4	142.3	0.540
Sri Lanka	80.18	75.06	182.7	177.2	0.426
Malaysia	3.675	2.479	6.891	4.904	-0.016
Oman	0.3393	0.292	0.6909	0.617	0.134
Philippines	39.72	31.49	100.2	69.41	0.106
Senegal	604.3	439.9	959.0	756.2	-0.009
Slovenia	287.2	252.4	350.3	276.1	-0.019
South Africa	7.735	8.132	11.56	11.64	0.206
Zambia	5235.3	4111.0	8115.5	7789.5	0.395

Notes: PPPs are calculated using consumption data from ring countries only; the aggregation uses a two-stage CPD and weighted-CPD method below and at BH level. The UK is the base country, and all PPPs and exchange rates take GBP as numeraire. The inflation rate is the change in the log of each country's CPI from the World Development Indicators minus the change in the log of the British CPI, so that the column shows the inflation rate relative to Britain. The exchange rates and PPPs for Estonia and Slovenia in 2011 have been multiplied by 15.6466 and 239.64 respectively, which are the rates at which each country entered the Eurozone. If extrapolation of PPPs worked perfectly, the log of the second column minus the log of the third column should be equal to the last column.

Table 3: Changes in log PPPs compared with changes in log CPIs for ring countries

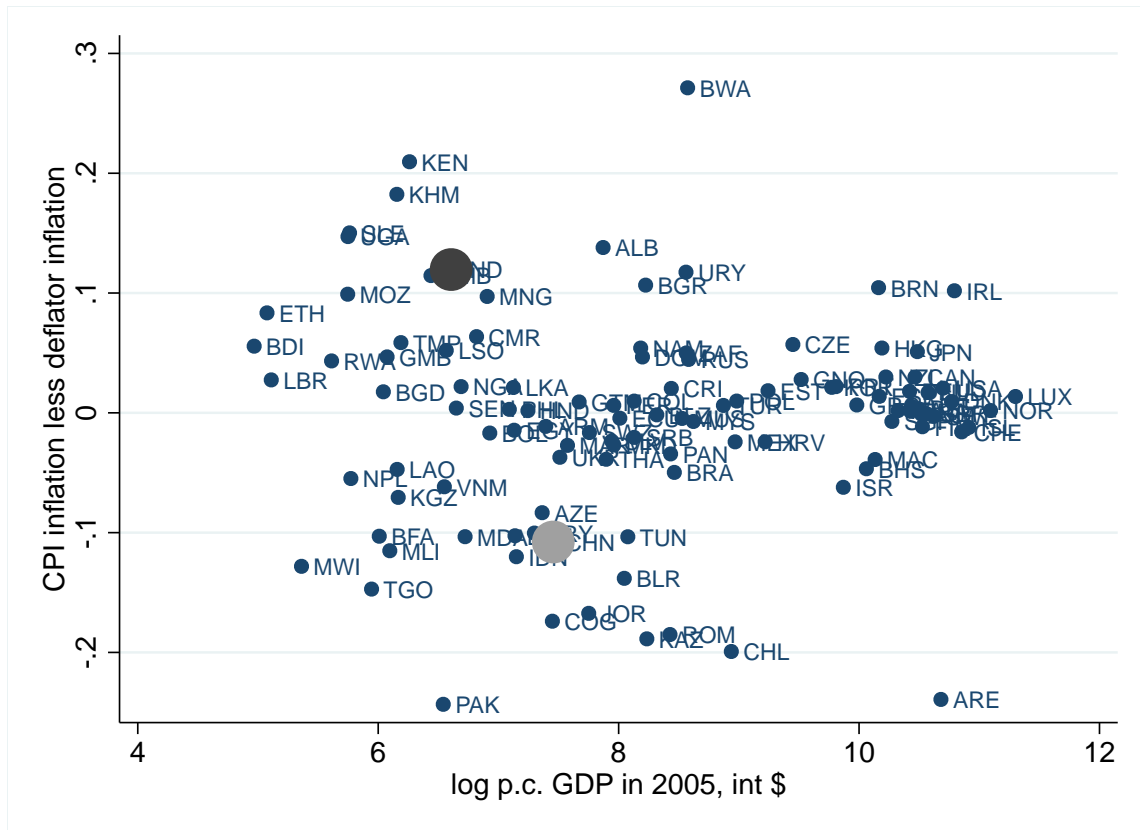
	Poor group	Rich group	Pooled
Intercept	-0.348 (9.2)	-0.090 (3.0)	-0.349 (10.0)
dln(CPI)	0.557 (4.4)	0.584 (2.4)	0.563 (5.2)
Poor group dummy	--	--	-0.260 (5.9)
# observations	10	8	18
Distance between lines	0.26		0.26

Note: The “rich” group consists of Britain, Japan, Chile, Hong Kong, Brazil, South Africa, Estonia, and Slovenia. The “poor” group consists of Zambia, Kenya, Cameroon, Senegal, Egypt, Sri Lanka, Jordan, Philippines, Oman, and Malaysia. The first two columns are separate regressions for the two groups, and the distance in the last row is the average distance between the lines. The last column is a regression for all 18 countries with the distance estimated as the intercept.

Table 4: Balassa-Samuelson regressions for Africa, Asia and Western Asia, 2005 and 2011

	US as reference			Eurostat-OECD GM as reference		
	2005	2011	Pooled	2005	2011	Pooled
Intercept	-0.00 (0.0)	-0.42 (4.4)	-0.09 (1.5)	-0.17 (9.6)	-0.51 (10.2)	-0.231 (5.0)
Log cons p.c.	0.170 (9.6)	0.110 (4.4)	0.146 (9.6)	0.170 (9.6)	0.110 (4.4)	0.146 (9.6)
2011	--	--	0.238 (6.3)			0.215 (5.7)
# obs	72	79	151	72	79	151
Distance	0.205		0.238	0.202		0.215
θ estimate	0.180		0.209	0.173		0.188

Note: The estimate of θ is calculated from dividing the mean distance between the two regressions by one plus the coefficient of log relative p.c. consumption, or in columns 1 and 2, by one plus the average of the two coefficients. The first two columns are separate regressions for the two groups, and the distance in the last row is the average distance between the lines. The last column is a regression for all 18 countries with the distance estimated as the intercept. Similarly for the right hand half of the table where the reference level of consumption is the geometric mean consumption per capita of the 46 countries that were in Eurostat-OECD in both rounds.



**Figure 3: CPI inflation minus consumption deflator inflation 2005-2011
India (dark grey) and China (light grey) enlarged**

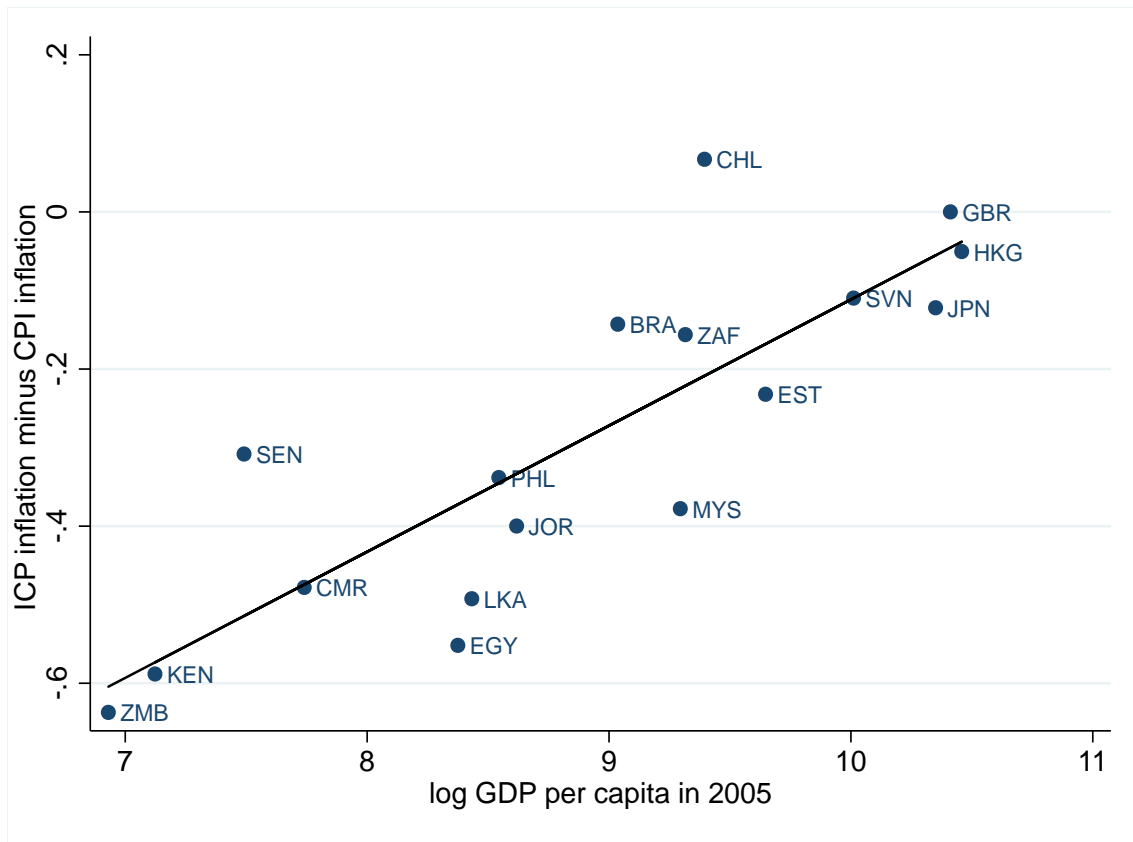


Figure 4: ICP minus CPI inflation and GDP per capita

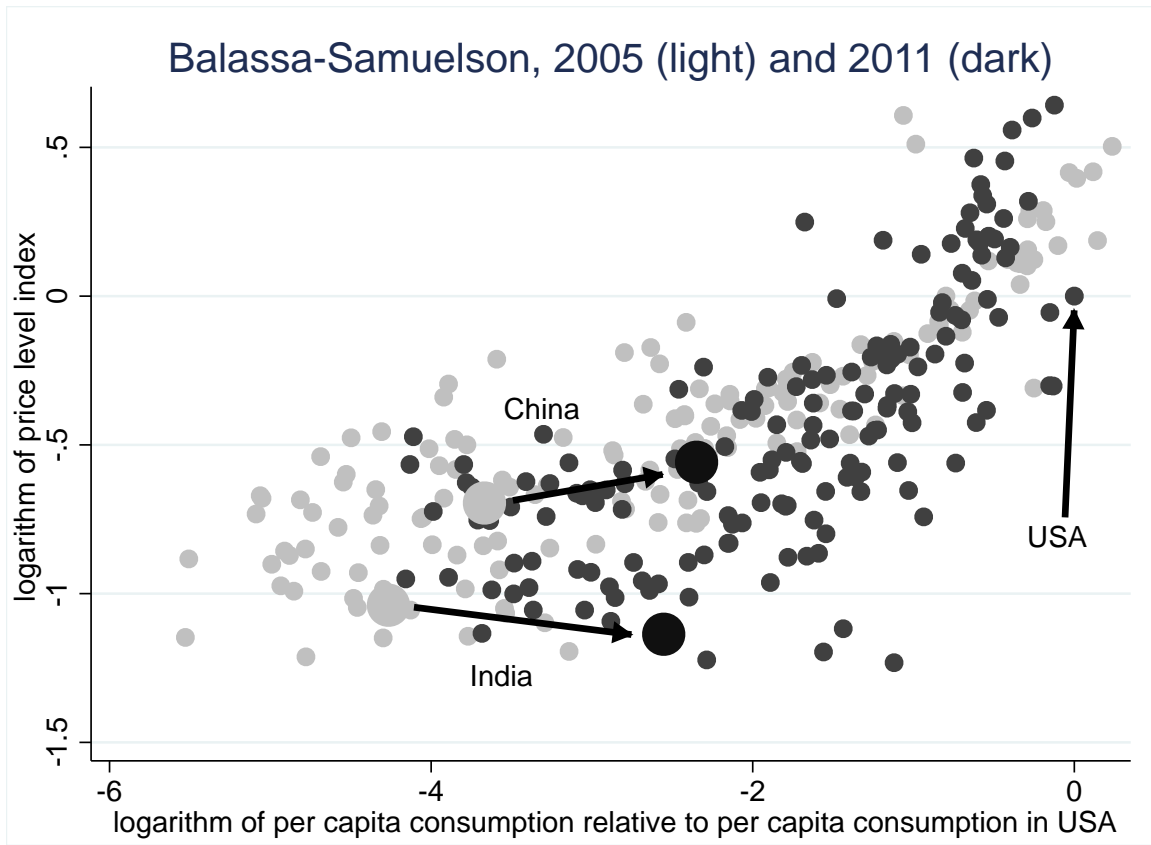


Figure 5: PLIs and consumption, 2005 (light) and 2011 (dark) ICPs