

Introduction:

Reshaping the World

ANGUS S. DEATON

The rounds of the International Comparison Program (ICP) are like successive Olympic Games. Similar to the Olympics, they do not happen every year, and in the first modern games only a few countries sent competitors, there were only a few events, and the standards of competition were relatively low. The participants were amateurs with day jobs, and, although they were great natural athletes, they did not take their training very seriously. Yet the first modern Olympics was a watershed, which eventually grew into the record-breaking professional event it is today in which almost all nations of the world come together in a truly global competition.

The ICP began in the late 1960s and early 1970s, led by Irving Kravis, Alan Heston, and Robert Summers from the University of Pennsylvania and Zoltan Kennessy from the United Nations. Like the Olympics, only a few countries (six) took part in the first round in 1967—four more were added in 1970—and prices were collected for only a small range of goods and services. Since then, each round has become bigger and better (and more expensive), with more countries represented, with more and more professional statisticians and economists involved, and with lots of preparatory training in the form of expert workshops, theoretical papers, and figuring out how to deal with problems that could not be solved in the previous round. The 2005 round of the ICP was by far the most professional, the biggest, the most thoroughly researched, and the most international—with 146 countries. It was the first round to be organized by a Global Office housed in the World Bank. Its findings changed the economic map of the world.

The 2005 ICP revealed a world that was much more unequal than we economists and others had thought. It was not quite like discovering water on the moon perhaps, but it was like discovering that the craters were deeper or that the planets were farther from the sun than we had always thought. And when the World Bank reworked the global poverty counts using the new data, it also found a world that was much poorer than it had previously thought.

The gaps between rich countries and poor countries—which we long knew were enormous—were even larger than previously measured. The average gap in the per capita gross domestic product

(GDP) in 2005 between two randomly selected countries in the world was about 5 percent larger as a result of the new data. For some individual pairs of countries, particularly a pair in which one was rich and the other was poor, the reshaping was much larger. The ratio of China's per capita income to U.S. per capita income was 40 percent smaller than it was based on earlier data. Much the same was true for India. And for many of the countries in Sub-Saharan Africa the widening of the gap was larger still. Meanwhile, what was true for countries was also true of individuals, and the average difference between the rich and the poor of the world was newly enlarged. As a consequence, the world had many more poor people below any global poverty line fixed in rich country currency, although, as will be seen, this is not the only way of setting the line.

Comparing Countries

What is the ICP good for? Why do we need it? And how did the world manage before it began? When it works well and the ideas match the measures, the ICP allows us to make sound comparisons of living standards between countries and between widely separated periods of time. The ICP collects the prices of thousands of items in each country and averages them to calculate price indexes for GDP, for consumption, and for its components. These indexes allow us to make international comparisons of the price of rice, or the price of food, or the price of all consumption items. The national accounts of each country reveal how much its citizens spend on rice, on food, or on all consumption, so that the price indexes from the ICP allow us to convert these money amounts, measured in local currency units, to “real” amounts expressed in a common unit, which is nearly always the U.S. dollar. The dollar amounts, such as Kenya's per capita GDP in U.S. dollars, is per capita GDP in Kenyan shillings (calculated by Kenya's statistical office) divided by the price index of Kenya's GDP in shillings per dollar.

These comparisons in common units reveal the relative sizes of different economies. They indicate not just that one country is richer than another, but by how much. Without the price indexes, it is impossible to calculate differences in living standards between countries or people's well-being in different countries, or to measure global inequality. Without them, it is also impossible to convert a global poverty line into its local equivalent, which is the number needed to calculate the number of globally poor in each country and therefore in the world. The World Bank's global poverty line is constructed from an average of the poverty lines of the world's poorest countries, and these local lines must be converted into international dollars before they can be compared and averaged.

Since World War II, a uniform set of principles for measuring national income has been in place. The principles evolved by Richard Stone, James Meade, and Maynard Keynes in wartime Britain were codified under UN auspices after the war under the guidance of Stone. These principles have since evolved into successive versions of the UN's *System of National Accounts*, or SNA, the latest in 2008 (Commission of the European Communities et al. 2008). In following this system, each country provides estimates of national income in its local currency, and this process, at least in principle, is carried out in the same way everywhere.

When trying to compare economic characteristics across countries, the obvious method is to use market exchange rates to convert everything into a common currency—such as the U.S. dollar—but conversion using exchange rates does not do a very good job. Many factors—such as movements of speculative capital—affect the exchange rate in the short run, so that the rupee-to-dollar exchange rate may fluctuate from day to day, even though neither India's nor the United States' living standards are changing. Expectations about the future can affect current exchange rates—for example, between

the euro and the dollar—even though there is no change in the current levels of income in Europe or the United States.

If all goods and services were freely traded between countries, traders would iron out these fluctuations, at least in the long run. But many goods and services are not traded at all—such as housing, many government and private services, the law courts, police services, haircutting, waiting tables, or babysitting—and there is no way in which to bring the prices of these items into line. In poorer countries, where labor is cheap, these nonexportable goods and services tend to be relatively cheaper than traded goods (such as wheat, gasoline, cameras, or machine tools), so that if common international units are used to value these nontraded goods, poor countries look less poor relative to rich countries than if domestic prices converted at market exchange rates are used.

All of this is just what every traveler knows. If an American gets off a plane in Delhi or an Italian disembarks in Addis Ababa and changes dollars into rupees or euros into birr, the amount of local currency received will go much further than the original dollar in Washington or the euro in Rome. In effect, the price level in poorer countries is lower than in richer countries. People in Delhi and Addis Ababa are indeed poorer than Americans, but because of the lower price levels they face, the difference is not nearly as large as it appears to be at market exchange rates. The alternative exchange rate that converts dollars and euros into rupees and birr in a way that preserves comparable purchasing power is called the purchasing power parity (PPP) exchange rate, and it is these PPPs that are measured by the ICP. In essence, PPPs are the price indexes computed from the hundreds of thousands of prices collected by the ICP.

The differences between market and PPP exchange rates are large and important. For poor countries, GDP per capita at international prices can be three (India) or four (Ethiopia) times larger than GDP per capita in domestic prices converted at exchange rates. But the ratio of market exchange rates to purchasing power parity exchange rates is not constant over time, nor is it the same for all countries with the same level of per capita income. So there is no choice but to actually collect the prices, and to do so, if not every year, at least on a regular basis.

Key Findings: Inequality

How did the 2005 ICP reshape the view of the world? The headline numbers came from India and China, whose economies “shrank” under the new estimates. The international dollar value of China’s per capita GDP in 2005 fell from \$6,757 in the 2007 *World Development Indicators* (WDI) to \$4,088 in the 2008 WDI (World Bank 2007, 2008). For India, the same comparison shows a reduction from \$3,453 to \$2,222. All of these numbers are for a single year, 2005, and because they come from converting the same local currency values but at different PPPs, another way of stating the change is that the PPP for China rose by a factor of 1.65, while the PPP for India rose by a factor of 1.55. Recall that GDP in international dollars is obtained by dividing a country’s own GDP by the PPP measured by the ICP, so that higher PPPs translate into lower estimates of GDP. The reduction in China’s and India’s GDP stems from the fact that the price index for China relative to that of the United States was 1.65 times higher than previously estimated, and that for India relative to that for the United States was 1.55 times higher.

Because international comparisons are carried out in international dollars, and because everyone is familiar with U.S. dollars, the obvious first interpretation of these data is that China’s and India’s economies are smaller than previously thought. But if the ICP had used not the U.S. dollar but, say, the Indian rupee as its unit of account, the change would have been that the U.S. economy was *much larger* than previously thought and China’s economy *slightly smaller* than

previously thought. All of these international comparisons are essentially relative; the ICP does not measure quantities, so it cannot say whether the absolute values of China's or India's per capita incomes were previously overestimated.

All of this may seem like hairsplitting, but it points to an important fact: the ICP widened the gaps between both India and China and the United States. Neither India nor China is any smaller or poorer (or indeed richer) than it was, although both are estimated to be smaller and poorer relative to the United States. In the 2007 *World Development Indicators*, the per capita income in the United States in 2005 was more than six times the per capita income in China, and more than 10 times the per capita income in India (World Bank 2007). In light of the 2005 ICP as reported in the 2008 *World Development Indicators*, these ratios increased to 12 times and nearly 19 times.

India and China are only two of the countries that moved farther apart from the United States in the 2005 ICP. Indeed, the effect was quite widespread, with many of the world's poorest countries shrinking relative to the United States. There was relatively little change among the world's richest countries (because many of them calculate PPPs every year, there is little opportunity for revision), so that the 2005 ICP caused a general widening of the dispersion of per capita incomes around the world.

Figure 1 plots the ratios of the "old" PPPs to the "new" PPPs against the logarithm of per capita GDP. Each point is a country, and the ratio is the ratio of the PPP reported in the 2007 *World Development Indicators* to the PPP reported in the 2008 WDI (World Bank 2007, 2008). If the ratio is greater than 1, measured per capita income has decreased relative to that of the United States; if it is less than 1, per capita income has increased relative to that of the United States.

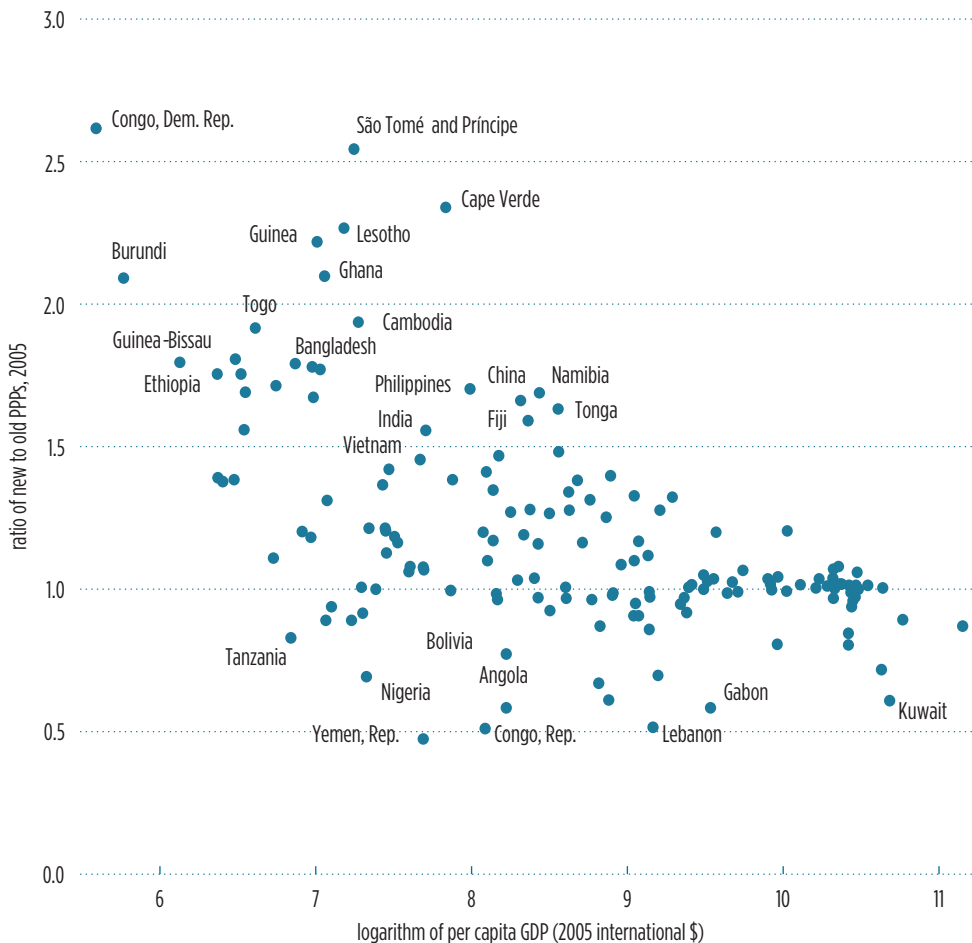
Figure 1 shows a strong downward slope, which means that the revisions of the 2005 PPPs were generally larger for poorer countries. As a consequence, many of the poorer countries are poorer relative to the United States, while the richer countries stay about where they were. Inequality between countries is therefore larger under the 2005 ICP. The upward revaluation of the PPPs for India and China turns out to be quite common, with many other countries in Africa and some in Asia experiencing similar or larger upward revisions. Indeed, the top left of the figure shows that some African countries had much larger upward revisions than India and China. A number of these had never been benchmarked in an ICP, and so the previous PPPs were little more than imputations or educated guesses.

Branko Milanović (2009) has calculated the Gini coefficient for income inequality among all the citizens of the world. This number is much bigger than the Gini coefficients for even the most unequal of individual countries because world inequality is dominated by differences *between* countries rather than by differences *within* them. According to Milanović's calculations, the world Gini coefficient in 2002 rose about 5 percentage points because of the revisions in the 2005 ICP, from 66 percent to 71 percent. Even if we ignore inequality within countries and compute the world Gini coefficient on the (counterfactual) assumption that everyone in each country has the same income, there is a similar increase of 5 to 6 percentage points just from the ICP revision.

Key Findings: Poverty

If the ICP made the poor world poorer relative to the United States, did it increase global poverty? Not necessarily, because the outcome depends on whether poverty is viewed from a rich country perspective or from a poor country perspective.

From a rich country perspective, the global poverty line is taken to be a dollar a day and is held fixed in real dollars. The global line in use before the 2005 revision was not precisely a dollar,

FIGURE 1 Ratios of Old to New PPPs in Relation to Income

Source: ICP 2005.

but \$1.08 in 1993 prices. By 2005 consumer prices in the United States had risen by 35 percent, and so the dollar-a-day line was actually \$1.46 in 2005 prices. When that global line was used *with the old PPPs* to calculate global poverty, the global poverty count was 931 million people. If the same global line, \$1.46 at 2005 prices, is used *with the new PPPs*, the global poverty count increases to 1.76 billion people, almost twice as many as before. Because the global poverty line is fixed in U.S. dollars and because the PPPs of poor countries have increased, the local equivalents of the global line have increased, and many more poor people are beneath them. Relative to the United States, then, the poor world is poorer than was thought, and there are many more poor people.

But use of this rich country perspective is not the only way to make the calculation. In 2005, at the old PPP of 11.3 rupees to the dollar, \$1.46 was worth 16.5 rupees in India (this figure lies between India's two domestic poverty lines of 17.7 rupees for urban India and 12.0 rupees for rural India). Thus using the old PPPs, at a global poverty line of 16.5 rupees per person per day, there were 931 million poor people in the world, which is just a restatement of the old dollar calculation. However, if the global poverty line is fixed not in dollars at \$1.46 but in rupees at 16.5, and if the

new PPP exchange rates are used, the new global poverty estimate is 943 million people, which is close to the original number. Relative to India, then, the world is neither poorer nor richer than first thought; the ICP revision has had very little effect.

One feature of this second calculation is worth noting. At the new, higher PPP for India, the global poverty line of 16.5 rupees is now worth only \$1.04 in 2005 U.S. dollars. This number is not only lower than the global poverty line in 2005 dollars (which it must be because of the increase in the PPP), but also actually lower than the global poverty line in 1993 dollars! However, this is just a consequence of the happenstance that the proportional increase in India's PPP was larger than the U.S. rate of inflation from 1993 to 2005.

Which of these two approaches is right, and why did the World Bank arrive at yet a third answer? Taking the second question first, the World Bank uses (a version of) the poor country perspective and calculates its global poverty line not as India's line but as the average of the poverty lines of the world's poorest countries. In principle, this approach should yield something like the just-described India-based calculation. And yet the Bank's calculations using the new PPPs show that 1.37 billion people are living in poverty, a substantial increase over the original estimate. The Bank arrived at this figure because it made other changes while implementing the ICP revision. In particular, it took the opportunity to update the group of countries whose poverty lines were used to calculate the global poverty line, and it turned out that, on average, the new reference group had higher poverty lines than the old reference group. Much of this was attributable to one country alone. India, which has one of the lowest poverty lines in the world but is no longer one of the poorest countries in the world, was dropped from the group, so the global poverty line went up. As the India example shows, it was this change in the global poverty line, not the ICP revision, that was responsible for increasing the global poverty count.

Of course, there is no *right* answer here. A good case could be made for holding the line fixed in dollars: the international community understands rich world currencies, is justifiably appalled by the number of people living on an unimaginably small but comprehensible amount, and is confused by a standard that appears to be denominated in dollars but is actually denominated in poor country currencies. The case in favor of the poor country standard is based on the reasonable belief that the poverty lines of the world's poorest countries are likely to be a good indication of the absolute minimal standard of living anywhere in the world. (But note that the Bank's new poverty line of \$1.25 at 2005 prices is substantially *above* India's rural poverty line, beneath which live nearly a quarter of a billion people.) Likewise, there is certainly a case for revising the line from time to time, and there is no compelling reason to always use the poverty lines of the same set of countries. Even so, the combination of a revision of the line and a revision of the PPPs at the same time is certainly confusing, and has made the whole process—which has always been hard to explain—even less transparent than usual.

More Countries in 2005, Fewer Imputations

What did the 2005 round of the ICP do differently? Do these changes make the new results more or less credible than the earlier numbers?

The most obvious improvement in 2005 was the increase in the number of participating countries. The 2005 ICP collected prices for 146 countries in all regions of the world, including 48 countries in Africa, a continent that is often underrepresented in international statistical compilations. Meanwhile, China was a full participant for the first time, and India participated for the first time since 1985. The only major gaps in 2005 were in Central America, the Caribbean, and a number of small island economies.

The very first ICP, run jointly by the University of Pennsylvania and the United Nations Statistical Office, collected prices in only six economies: Hungary, India, Japan, Kenya, the United Kingdom, and the United States. Four other countries—Colombia, France, Germany, and Italy—collected data for 1970, and were included in the first ICP set of PPPs published in 1975. With successive rounds, the number of countries gradually increased, reaching 60 in the 1980 round, 64 in the 1985 round, and 118 in the 1993 round, the last before the 2005 round.

For most academic economists who use them, the results of the ICP are accessed through the Penn World Table (PWT), Mark 1 of which appeared in 1980. Mark 5 (Summers and Heston 1991), based on the 1985 round, contained results for 139 countries and covered the period from 1950 through 1988, though not with all years for all countries. Mark 5, described by Robert Summers and Alan Heston in an important article in the *Quarterly Journal of Economics* in 1991, was responsible for reigniting an academic interest in the empirical study of economic growth, and there is now a huge literature using these data as well as the later versions of PWT Mark 6 based on the 1993 round. Mark 7, using the 2005 ICP, is currently in preparation.

For countries not covered by the ICP, PPP exchange rates are “filled in” by estimating the price level based on each country’s level of development. For example, in the examples just cited, the price level for India is 0.33 and for Ethiopia 0.25, so that for a country with per capita incomes between the two the price level would be somewhere around 0.30, and the PPP would be 30 percent of the market exchange rate. In practice, prediction of the price level can be improved by taking into account factors in addition to per capita GDP. However, each country is special in some way, and the predictions of a regression are never as good as using actual data.

In past rounds, when a country missed an ICP round, such as India in 1993, a guess could also be made by taking a previous PPP exchange rate, from the 1985 benchmark, and “updating” it from the relative rates of inflation in the United States and India. But the basket of goods in each country’s consumer price index (CPI) is not the same as the international baskets used in the ICP. Nor are domestic CPIs always constructed to the same principles. So this, too, is only a rough and ready substitute for collecting the data.

In summary, one of the great strengths of the 2005 ICP was that very few imputations and updates were required because it covered all major countries together for the first time.

Better Linking of the World in 2005

An important improvement in the 2005 ICP was the way in which the price collection was organized. In the early days when only a few countries were participating, the ICP was carried out centrally—for example, at the University of Pennsylvania—but as the number of countries grew over time, price collection was regionally dispersed. Each region calculated its own set of regional PPPs relative to a regional base country, with PPPs for the world calculated at a final “linking” stage.

By the 1993 ICP, the dispersal had reached the point at which the central organization had become very weak. This situation caused many problems, and one of the main aims of the 2005 round was to develop a coherent global structure for the ICP. It was at this point, too, that the World Bank was brought into the ICP and became not only a major funder but also the home of the Global Office, which was responsible for the overall design of the project and for combining the regional estimates into a set of global PPPs. Each of the five regions had its own office, each maintained its own data collection machinery, and each calculated its own set of PPPs for the region. Meanwhile, Eurostat and the Organisation for Economic Co-operation and Development (OECD) were jointly conducting their own regular process of calculating PPP exchange rates for their countries, but in

coordination with the ICP regions. For other regions, such as Africa, the 2005 ICP was a new regional effort with an emphasis on statistical capacity building that would not have taken place otherwise. At the center of the regions was the Global Office in the World Bank, which was responsible for coordination and for the final linking stage in which a global set of purchasing power parity exchange rates was calculated from the information submitted by the regions. The Global Office worked under the auspices of an executive board formed by the UN Statistical Commission and was provided with technical advice by the panel of experts who formed the Technical Advisory Group. Overall, the World Bank furnished the organizational and technical skills to make this enormous operation work.

What were the payoffs from this reorganization? What had gone wrong in the 1993 round, and how did the new structure help to remedy it?

The 1993 round was neither centrally coordinated nor controlled, and in the face of underfunding at the center it became a set of regional exercises carried out at different times, each of which collected data and calculated regional PPPs. A United Nations report circulated in 1998, commissioned jointly by the UN, the World Bank, and the International Monetary Fund, and commonly referred to (after its chairman) as the Ryten report, argued that the estimates from the 1993 ICP were not credible. It concluded, with faint praise, that the “ICP is a programme worth keeping but that its current condition, if little is done about it in terms of credibility, quality of output, and survival prospects, is poor” (United Nations 1998).

One credibility problem came from the way the global PPPs were constructed. Without adequate central coordination, not all of the planned links could be carried out, so that the global PPPs were calculated by linking the regions *ex post*. The linking was accomplished by using countries included in more than one region as bridges. This approach is conceptually similar to linking an old and a new time series from a price index through its value in a bridge year for which both price indexes are available. But spatial price indexes cover many countries simultaneously and do not have the natural ordering that comes in time series. This difference means that the results of linking two regions through a common country will depend on which country is used, a choice that needs to be made on principle and not by happenstance, as was the case in 1993. One particular concern is whether the linking country is special in some way—for example, whether it has patterns of consumption and relative prices that are somehow unusual, something that is often thought to be the case for Japan, which was one of the linking countries in 1993.

The results also depend on just how the linking is done. For example, one possibility is to use the PPP exchange rates between Japan and India, both in the Asia-Pacific region, and Japan and the United States, both in the OECD region, to derive a PPP exchange rate between the United States and India. A more detailed exercise can be done to convert the price of individual goods and services in India—rice, clothing, automobiles—from rupees into dollars using the price of each good in Japan as a bridge. As with the choice of bridge countries, the level of disaggregation will affect the final answer. The spirit of the ICP dictates that the linking be done at the finest level possible, but without central coordination this, too, was dictated by happenstance, and from the uncoordinated choices of each region.

China actively participated in 1993, at least to the extent of making a number of comparisons between cities in China and elsewhere, but those were never fully incorporated into the 1993 ICP. As a result, the PPPs for China in 1993 were based on data collected in 1986 to make a bilateral comparison with the United States, and then extrapolated forward to 1993.

In the 2005 ICP, the linking of the regions was centrally planned and implemented. Instead of relying on a few countries that happened to be in more than one region, a group of 18 “Ring” countries¹ was selected in advance, with two or more countries in each region. Each Ring country carried out a second round of price collection, relying on a common special list of more than 1,000 items. The Ring can be thought of as a separate, self-contained mini-ICP, although “mini” is

relative, because the number of countries in the Ring was larger than the number of countries in the first round of the ICP. It generated a set of Ring PPPs, and beyond that a set of Ring prices—in common Ring currency—for each of the goods and services in the ICP. These prices, appropriately averaged, were then used to “glue” the regions together into a global set of accounts in which there was a purchasing power parity exchange rate for each country (relative to the U.S. dollar, which was the numeraire), but also a set of prices—in U.S. dollars—for each of the 155 goods and services (“basic headings”) covered by the ICP.

The linking of the regions in the ICP 2005 was not without its problems. The most serious of these were not failings of the ICP itself, but came from the conceptual difficulties associated with all exercises of this kind, particularly when comparing countries whose patterns of consumption and relative prices are radically different. It is one thing to make PPP comparisons of France and Germany, or of Kenya and Tanzania, but it is on altogether shakier ground to compare Canada with Cameroon, Japan with Senegal, or Bolivia with Tajikistan. Such comparisons are difficult in theory and subject to a wide margin of uncertainty in practice. Indeed, this is something that anyone using the results of the ICP should always keep in mind.

That the linking procedures in the 2005 ICP were well documented and well thought out, even if not unchallengeable and certainly not the final word, made the 2005 round much more credible, more reliable, and safer than any of its predecessors.

More Precise Definitions of Goods and Services in 2005

With its better coordination, the Global Office was able in the 2005 ICP to provide the technical support needed to help each country collect prices in a coherent way and to check and edit the results for credibility and correctness. Such advances attract little attention from the outside, and any description is soporific both to write and to read, but their importance is hard to exaggerate. One of the criticisms in the 1998 UN report was that the ICP was very strong on the theory of the index numbers underlying the PPPs, but much weaker on giving statistical offices precise instructions on how to collect prices. That weakness was remedied in 2005.

In the 2005 ICP, each region developed its own list of prices—something that makes sense when countries are more similar within regions than across them. The Ring list was put together centrally by the Global Office, based on inputs from each region. This list is crucial in linking the world and plays a central role in determining the distance between poor and rich countries, and the extent of world inequality. The 2005 Ring included countries as disparate as Senegal and Cameroon in Africa, Japan and Estonia in the OECD, Jordan and Oman in Western Asia, and Malaysia and the Philippines in the Asia-Pacific.

Any list that runs across such countries has to satisfy two criteria that are often at odds. One criterion, in order to validate the international comparisons, is that the goods being priced are the same in all countries. This criterion calls for precise definitions of the goods in the list. If the definitions are too vague—for example, a “shirt” or a “family car”—the ICP runs the risk of pricing lower-quality items in the poorer countries, so that it is not comparing like with like and is underestimating (overestimating) price levels in poorer (richer) places.

The second criterion is that the goods in the list be widely consumed in each country, so that the goods are genuinely representative of what people buy.

The 1998 UN report noted the difficulty of satisfying both of these criteria at the same time, as well as the consequences for the credibility of the ICP, but it did not propose any solution. In any

event, the 2005 ICP dealt with the issue by developing very precise lists, especially for the Ring list. For example, instead of wine, or even red wine, the item was “Bordeaux *supérieur*, with state certification of origin and quality, alcohol content 11–13 percent, vintage 2003 or 2004, with region and wine farmer listed.” This level of detail clearly does very well according to the first criterion of pricing the same item everywhere. The second criterion was dealt with by asking enumerators to report whether the item was representative of local consumption, with the aim of down-weighting non-representative products. For a number of reasons, including the difficulty of defining *representative*, the reporting did not work in some regions. Nevertheless, it was clear that the 2005 ICP was a huge improvement on one of the two criteria, if not on both. Making both criteria work remains an active research area, and there will undoubtedly be further refinements in the 2011 ICP.

Continuing Progress on Other Issues

The ICP has long had a list of problem children known as “comparison-resistant” goods and services. Many of these are services for which it is traditionally difficult to define quantities—for example, how does one compare a hip replacement or brain surgery in Nairobi, Tokyo, and Buenos Aires?—and many relate to government activities—that is, the provision of education, defense, or administrative services by civil servants. These are all areas in which there are long-standing problems of measurement, even for domestic national accounts, and these problems tend to be even more difficult in cross-country comparisons. The handling of these issues occupies a large portion of the time of the technical committees that support the ICP. None of them is definitively solved, and none of the current solutions is above criticism. But there is also no doubt that progress is being made, and that better methods and better data collection are constantly being brought on line.

The ICP relies heavily on data it does *not* collect—the national accounts of the participating countries. Because the ICP collects data on *prices*, not on expenditures or quantities, when it reports levels of real income in international dollars in different countries, it is relying on local estimates of income in local currency and then converting them to international dollars by dividing by the price indexes from the ICP. Indeed, even the construction of the price indexes from the prices of individual goods and services relies on the local national accounts to provide the weights that indicate the relative importance of each category. Thus the ICP comparisons are only as good as the national accounts that go into them, over which the ICP has no direct control.

Put more positively, the ICP itself is an opportunity for the Global Office to help countries improve their national accounts. A good deal of this was done in 2005, and more is being done now in preparation for the 2011 ICP. Thus the ICP, like the Olympic Games, can leave a lasting legacy of better local infrastructure.

Credibility of the ICP Revisions

According to figure 1, the 2005 ICP made few revisions among the richer countries, and there were essentially no revisions among the Eurostat-OECD countries. These countries have their own PPP program, run by the European Union and OECD statistical offices, which calculates PPPs on an annual basis, and which was incorporated into the 2005 ICP. For those countries with annual monitoring, there is no possibility of the large revisions that can happen when a country has not been benchmarked for a dozen years or more. The large number of rich countries without revision illustrates the benefits of calculating PPPs at a higher frequency than is the case for much of the rest of the world.

Are the revisions elsewhere credible? Are the new PPPs more reliable than the old ones? Are the higher inequality measures better than the old ones?

The answer to all of these questions is certainly affirmative. As documented in the UN report, the 1993 ICP was in some disarray and had lost much of its credibility. This disarray was particularly evident in the way the regions were linked, and it is the linking that is responsible for establishing the PPPs in Africa and Asia relative to the United States and the other OECD countries. The linking in the 2005 ICP through the Ring countries was well thought out in advance and centrally and systematically implemented.

In addition to the improvements in the linking itself, there were many more countries to be linked. Most of the African countries had never been benchmarked, nor had China, and the Indian benchmark was more than 20 years out of date. Even if there had been no linking and the 2005 ICP had produced only a set of unconnected regional accounts, it would have been a huge advance in the supply of credible price information from all countries.

Of course, it is always good to keep in mind that international comparisons are difficult, especially between countries that are very different in their consumption patterns and in the structure of relative prices. No matter how accurate, detailed, and careful is the price collection by the ICP, comparisons of, say, Senegal and Japan, Brazil and Bangladesh, or the United States and Tajikistan are going to be rough at best. Indeed, a good starting point for anyone using the ICP results is to take such comparisons with a large grain of salt.

One central issue also identified in the UN report was how to resolve the conflict between, on the one hand, wanting to measure the same goods in different places and, on the other, ensuring that the goods whose prices were being measured were representative of consumption in each country. In the 2005 ICP, this conflict was resolved in favor of ensuring that the items were closely comparable. If the list of such items contains many items that are common in rich countries but rare and expensive in poor countries, it is possible that linking through the Ring would exaggerate the difference in prices between poor and rich countries, and this exaggeration would contribute to an overstatement of global inequality.

However, work by this author (Deaton 2010) failed to yield much evidence of this effect in the details of the Ring comparison in the 2005 ICP, or at least that the effect contributed very much to the widening of the gap between rich and poor countries. Instead, the main source of uncertainty is a more fundamental one—because of the different relative prices and different weights, there is a wide range of reasonable ways of calculating PPPs. This issue cannot be resolved by better measurement, although research can certainly build on what has been done so far in order to suggest new measurement. In the meantime, transparency about methods and about data is of the greatest importance. To aid this, the Global Office has provided data sets to researchers that allow methodologies to be compared, and this analysis will surely guide further improvements in the 2011 ICP and beyond.

NOTE

1. This group included the economy of Hong Kong SAR, China.

REFERENCES

Commission of the European Communities, International Monetary Fund, Organisation for Economic Co-operation and Development, United Nations, and World Bank. 2008. *System of National Accounts 2008*. <http://unstats.un.org/unsd/nationalaccount/SNA2008.pdf>.

- Deaton, Angus. 2010. "Price Indexes, Inequality, and the Measurement of World Poverty." *American Economic Review* 100 (1): 5–34.
- Milanović, Branko. 2009. "Global Inequality Recalculated: The Effect of New PPP Estimates on Global Inequality." Policy Research Paper 5061, World Bank, Washington, DC, September.
- Summers, Robert, and Alan Heston. 1991. "The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950–1988." *Quarterly Journal of Economics* 106 (2): 327–68.
- United Nations. 1998. "Evaluation of the International Comparison Program." UN Economic and Social Council, E/CN.3/1999/8. <http://unstats.un.org/unsd/methods/icp/8-e.pdf>.
- World Bank. 2007. *World Development Indicators 2007*. Washington, DC: World Bank. <http://data.worldbank.org/data-catalog/world-development-indicators>.
- . 2008. *World Development Indicators 2008*. Washington, DC: World Bank. <http://data.worldbank.org/data-catalog/world-development-indicators>.