

SIX

Globalization and the Greatest Escape

IN THE YEARS SINCE World War II, the modern world has seen the greatest escape of all. Rapid economic growth in many countries has delivered hundreds of millions of people from destitution. Material wellbeing has risen as death rates have fallen, and people are living longer and richer lives. As always, progress has not been even; some of the most rapidly growing countries have narrowed the gap with the rich countries, but their progress has opened up new gaps between them and the countries left behind. Once-poor countries in Asia have moved into the middle, leaving chasms between them and many countries in Africa.

The fall in mortality, especially among children, caused world population to increase at a rate that has no precedent in human history, a true explosion of population. That global poverty should have fallen in the face of such an increase in numbers would have astonished most commentators of the 1960s, for whom the looming "population bomb" threatened living standards everywhere in the world. The great economist and Nobel laureate James Meade used to complain that the three great disasters of the twentieth century were the "infernal" combustion engine, the population explosion, and the

Nobel Prize in economics. On the population explosion, most of his contemporaries would have agreed, and even today many continue to see population growth as a serious threat (along with the infernal combustion engine). Yet not only has the world added four billion people over the past half-century, but the seven billion who are alive today have also, on average, much better lives than their parents and grandparents.

Averages are no consolation to those who have been left behind. We have already seen that the average growth in the United States was far from equally shared. The United States is not the only country where inequality has been rising, and although there are important exceptions, rising income inequality is a common recent experience in many countries. What about inequality *between* countries? Many once-poor countries have seized the "advantage of backwardness": the opportunity to adopt (and even improve) the knowledge and technologies that have long been known in the now-rich countries. The catch-up countries can bypass the long process of trial and error that limited growth in the past. Countries such as the Asian tigers—Hong Kong, Singapore, South Korea, and Taiwan—and more recently China and India, have experienced rates of economic growth that are multiples of anything seen previously. But growth has been unequally distributed, and most of the countries that were poor fifty years ago have not been able to emulate China, India, or the tigers.

Perhaps surprisingly, and in spite of the achievements of the fast growers, there has been little or no narrowing of income inequality between countries; for every country with a catch-up story there has been a country with a left-behind story. The spread of average incomes between poor countries and rich countries is as large as it has ever been. If we arrange countries by their average incomes, from the poorest to the richest, we can compare the country that is a quarter of the way up from the bottom—a moderately poor country—with the

country that is a quarter of the way down from the top—a moderately well off country. In 1960, the moderately well-off country had an average income that was seven times that of the moderately poor country; by 2009, that ratio had risen to eight and a half times.

This chapter looks at this postwar miracle—the greatest of great escapes: how it came about and how it closed old inequalities and opened new ones. We shall also look closely at the numbers and ask whether it is safe to believe them. Global measures of poverty and inequality are fraught with difficulties. We know less than we should and certainly less than one might think from reading and listening to the torrent of popular pronouncements.

Measuring the World

Measuring material wellbeing is no easy matter, and even *income*, an everyday term, is hard to nail down precisely. Our measures of poverty and inequality are only as good as our measures of income. Life gets even tougher when we want to make comparisons across countries. People have a pretty good idea what sort of income is needed to avoid being poor in the community in which they live. Even if national poverty lines do not capture what it costs to live in your community, let alone differences of opinion about needs, we can still hope that most citizens and policymakers will see the national poverty lines as reasonable figures that divide those who are getting by from those who are not. But if we are to count the poor of the whole world, we need a single poverty line that makes sense in Nairobi and Quito, Karachi and Timbuktu, and perhaps even in London and Canberra. For this, as for any international comparison of incomes, we need to be able to convert one currency into another, and it turns out that, for this purpose, foreign exchange rates are useless.

A good place to start, then, is by asking how to convert one currency into another, for example, dollars into rupees. There is an exchange rate, varying from day to day, which is the number of rupees that you can buy in the market for a dollar; as I write this, in April 2013, that rate is 54.33. So if I take a flight from New York to Delhi and go to the foreign exchange counter, I will get around 50 rupees for each dollar, or perhaps less depending on the bank's commission. Yet when I get into the city, I will find that, even in the most expensive hotel, I can buy a lot more for 50 rupees than I can buy for a dollar in New York; if I go to the canteen at the Delhi School of Economics, or buy street food, the difference is enormous.

One simple way of saying this is that the price level is lower in India than in New York; if money is converted at the market exchange rate, most things in India are cheap compared with their prices in the United States. In fact, according to the latest estimate, the price level in India is only about 40 percent of the price level in the United States, so that, if we take a typical bundle of things people buy, the bundle in India costs only 40 percent of what it costs in the United States. Put another way, prices would be the same in both places if the exchange rate were 20 rupees to the dollar, not 50 rupees. This “correct” exchange rate, the one that would make a dollar worth the same in both places, is called, appropriately enough, the purchasing power parity exchange rate or, for short, the PPP exchange rate. The PPP rate is the exchange rate from dollars into rupees that would give the same purchasing power in both places. If the price level is lower in Delhi than in New York, as it is in most poor countries, the PPP rate will be lower than the foreign exchange rate.

How do we know these numbers? There is no market in which currencies are converted at PPP rates, so there is no alternative but to go and find out what the prices are. Teams of international researchers

and statisticians collect millions of prices in countries all around the world and average them to get a price level for each country. The first such calculations were carried out for six countries in the mid-1970s by a team of economists at the University of Pennsylvania led by Irving Kravis, Robert Summers, and Alan Heston, the last of whom continues to work in the area and is one of those responsible for many of the numbers in this book. These pioneers changed the way that economists see and think about the world; without their work, we would have no idea how to compare living standards across countries.¹

One of the first things learned from these international comparisons is that my Indian example was and is quite common; price levels are lower in poorer countries, and the poorer the country, the lower is its price level. Many people find this conclusion both unlikely and surprising. How can we have a world in which things are cheaper in one place than in another? If something like steel or gasoline were much cheaper in Delhi than in New York, why would a trader not buy it in Delhi and get rich by selling it in New York? In fact, the prices of steel and gasoline are not very different once we take into account the costs of shipping and of local taxes and subsidies. But these arguments do not apply to everything. The fact that a haircut in Delhi or a dinner in Bangkok would be a splendid bargain in New York holds no attraction for a trader because these services are in Delhi and Bangkok, not New York, and it is not possible to get them from one place to the other. Because people are poor in poor countries, services are cheap there, but many of those services are immobile.

If everyone were free to migrate from one country to another, wages in rich countries would fall and wages in poor countries would rise, and the world would be a much more equal place. Of course, opposition to lower wages in rich countries is precisely why people are *not* permitted to migrate at will, and it is why meals and haircuts are

so cheap in poor countries. The price of land, like the price of labor, cannot be arbitrated between rich countries and poor countries. Cheap housing in India or Africa cannot be brought up to American prices by simply moving the land across the ocean. The presence of cheap land and cheap labor in poor countries explains why price levels in poor countries are so much lower than in rich countries. The market sets the exchange rate to equalize the prices of steel, gasoline, automobiles, and computers—everything that can be and is part of international trade—but the price level depends on goods and services that cannot be traded. Because those are cheaper in poor places, the poorer the country, the lower are the average prices.

Because of lower prices in poorer countries, we get the wrong answer if we use market exchange rates to make conversions of the costs of living. Newspaper reports almost always get this wrong, and even economists sometimes forget it. In the spring of 2011, the Government of India (unwisely and ungenerously) argued before the Indian Supreme Court that 26 rupees a day was enough to avoid poverty, at least for those living outside the cities. In the uproar that followed, the Indian (and international) media noted that even the World Bank—not seen by most Indians as a benevolent institution—used a poverty line of \$1.25, which, at the exchange rate of 53 rupees to the dollar, was more than twice as generous as the government's line. But at a PPP rate of 20 rupees to the dollar, the World Bank line is 25 rupees, close to the line that the government had suggested. Even the *Financial Times* used the market exchange rate to make the conversion from rupees into dollars and noted that the government's line was only \$0.52 as opposed to the Bank's \$1.25. A more accurate number would have been \$1.30—a miserable pittance indeed, but almost three times the wrong number.

The United Nations Development Programme made this mistake for many years, opening itself to charges of deliberately exaggerating

the poverty in poor countries. Any time we read about living standards in poor countries—whether a wage rate, what it costs to visit a doctor, or the price of food or transportation—and make the obvious conversion using a market exchange rate, the results will be too small by a factor of two to three. Wages are certainly low in poor countries—just another way of saying these countries are poor—but nothing is served by exaggerating just how poor they are relative to the rich countries of the world.

When we compare living standards across the world, or calculate global poverty or inequality, PPP exchange rates are always the right ones to use. The phrase “across the world” is important here; when we look across people in the same country to calculate inequality—as for the United States in Chapter 5—we might reasonably choose *not* to adjust for differences in prices across places. While it is certainly cheaper to live in Kansas or Mississippi than it is to live in New York City, there are more amenities in the city. Indeed, if people are free to choose where to live, the higher prices in the big city are likely to be a reasonable guide to the value of those amenities. If this is so, we can compare incomes across space without adjusting for prices; people with higher incomes in Manhattan, New York, really are better off than people with lower incomes in Manhattan, Kansas. Things are different when we compare the United States with India, or France with Senegal, where free movement of people is not possible. Even if living in America brings more amenities than living in India—and I do not know whether that is so—there is no reason to suppose that the difference between American and Indian price levels reflects those differences in amenities. So when we compare Indian and American incomes to assess world inequality, we must adjust for prices using PPP exchange rates.

PPP comparisons are better than comparisons at market exchange rates, but they are far from perfect. Price levels are calculated by col-

lecting the prices of comparable items in different countries, like a kilo of rice or a haircut in Hanoi, London, or São Paulo. Not all items are easily priced, however. How do we price the dwelling that a poor family built for itself in a village, or a shack in an urban slum? There are often no rental markets for these properties, just as there are no rental markets for many kinds of housing in richer countries. As in the United States, it is hard to know how to value the services that are provided by governments to their people—items like Medicare—and it is even harder to do this in a consistent way for all countries, making sure that we are comparing like with like. For a substantial share of what people spend, there are no market prices and we have to make do with guesses that, while always sensible, may be wide of the mark. This does not mean that it is better to use regular exchange rates, which we know are wrong—only that the PPPs, which are the right numbers, are subject to unavoidable uncertainty.

Think for a moment about collecting prices of comparable items in different countries. Say we are pricing men's shirts. In the United States, a standard item might be a dress shirt from a well-known manufacturer, say a Brooks Brothers oxford cotton button-down shirt. If we are to compare this with a man's shirt in Bolivia, the Democratic Republic of the Congo, or the Philippines, we are faced with a choice between two equally unsatisfactory alternatives. The standard shirt in those countries is likely to be something much cheaper and of lower quality than the Brooks Brothers shirt, so that, if we price it, we are not comparing like with like, and we will be *understating* prices in the poor country relative to the rich country. The alternative is to hunt really hard for the Brooks Brothers shirt—perhaps in the fanciest store in the capital city—but this carries the opposite risk: that we find the shirt, but only in a very specialized and very expensive store that is used by a handful of wealthy patrons. If we follow this route, we *overstate* prices in the poor country, at least

for ordinary people. There is a constant tug-of-war between two conflicting aims: collecting prices only on items that are internationally comparable and collecting prices only on items that are representative of what people buy. In extreme cases, comparisons are effectively impossible if something that is important and widely used in one country does not exist in another. Teff is the basic grain staple in Ethiopia and is rarely used anywhere else; tofu is important in Indonesia but not in India; alcohol is not available in many Muslim countries.

Even when all the prices are available, people spend their money on different things and in different proportions in different countries. One example will be familiar to anyone—like me—who was brought up in Britain and who now lives elsewhere. One of the basic necessities of existence for Brits is a product called Marmite. This is a (very) salty yeast extract that is a by-product of brewing, originally discovered by Louis Pasteur, who in turn licensed it to a British beer manufacturer. In Britain, Marmite is cheap and widely consumed; it comes in large black pots. In the United States, where I now live, Marmite is available, but it is expensive and comes in very small black pots. Marmite is a well-defined and precisely comparable item that is easily priced in both the United States and Britain. But if we compare prices in the United States and Britain by calculating the relative costs in the two countries of the goods that the British buy, including lots of Marmite, we will find that the United States is a very expensive place. If we look at relative prices using the goods that Americans buy—which include items like graham crackers or bourbon that are rare and pricey in Britain—we will find that Britain is the expensive place.

Comparisons between rich countries like Britain and the United States are not in fact very sensitive to whether we use American goods or British goods as the basis for the comparison, but the Marmite

example illustrates a basic issue that affects all international price comparisons. Countries tend to buy a lot of things that are relatively cheap at home and fewer things that are relatively expensive at home, so that comparing the cost of living abroad using the home country's basket will tend to overstate the costs of living abroad. If we use the foreign country's basket, we tend to understate the relative costs. In practice, the statisticians tend to split the difference and take an average.

Splitting the difference is a sensible solution, but it doesn't make the problem go away, as we can see from thinking about what happened recently when the price statisticians compared prices in Britain with prices in the West African country of Cameroon. In Cameroon, as in much of Africa, air travel is very expensive, and there is very little of it; ordinary people do not fly. In Britain, air travel is cheap, and even relatively poor people fly abroad on vacation. Pricing what the British do at Cameroon's prices makes Cameroon seem like a very expensive place. Averaging helps, but whatever we do, the price of air travel has a substantial effect on Cameroon's PPP—the price level in Cameroon would be 2–3 percent cheaper if air travel were ignored—even though there is close to no air travel in Cameroon. It is one of the unfortunate facts of life that international comparisons depend on things like this that—at least in some contexts, such as measuring poverty—make little sense. Once again, the trouble here is that Cameroon and Britain, unlike the United States and Britain, are very different from one another.

The comparison of Cameroon and Britain is not a big deal; what is definitely a big deal is the comparison of the United States and China. According to the latest estimates from the World Bank, per capita GDP in China in 2011 was \$5,445 compared with \$48,112 for the United States, so that per capita income in the United States is nearly

nine times that in China. But those calculations are at market exchange rates, and they take no account of the fact that the price level in China is about two-thirds of that in the United States. If we convert at PPP exchange rates, per capita income in China is \$8,400, so that the ratio of per capita incomes at PPP, a much better indicator of relative living standards, is only 5.7 times, not 8.8 times. For those worried about the absolute size of the two economies—the military or the diplomatic corps, for whom a country's influence in the world depends on *total* resources—we have to multiply by 4.31, which is the ratio of the Chinese to the American population. In total then, the Chinese economy is three-quarters the size of the American economy. Given that China is growing much more rapidly than the United States—on which more below—we might expect China to overtake the United States in the not-too-distant future—in only six years if China's growth rate is 5 percentage points higher than the U.S. growth rate.

The numbers in the previous paragraph treat the PPP exchange rate as if we really knew it in the way that we know the market exchange rate. But once we recognize the “Marmite” or “air travel in Cameroon” problems, as well as the uncertainties about making comparisons that are both representative and comparable, we should ask whether the true PPP exchange rate might be rather higher or rather lower. In my work with Alan Heston, we calculated that once the Marmite problem is recognized—or more precisely once we recognize that we can average the two sets of prices using either the Chinese or the American bundle of goods—there is about a 25 percent margin of error on either side of the PPP.² So we can say only that Chinese per capita income in international dollars in 2011 was somewhere between 13 and 22 percent of per capita income in the United States. The aggregate Chinese economy is somewhere between 56 percent and 94 percent of the aggregate American economy. While it is convenient to

split the difference—if only because we do not want to work with a vast range of possibilities—the fact remains that splitting the difference is an arbitrary way of resolving a conceptual problem to which there is no fully satisfactory solution.

In the very special case of China, there are many other issues to which I cannot do justice here. Perhaps the most important is the long-running and still unsettled debate over whether the official Chinese growth rate is too high to be credible, as many scholars think, and if so, by how much it should be adjusted down.

I do not want to leave the impression that international comparisons are impossible, or always subject to very large margins of error. In 1949 my mentor at Cambridge University, Richard Stone, asked, “Why do we want to compare the United States with, say, India or China? What possible interest is there in it? Everybody knows that one country is, in economic terms, very rich, and another country very poor; does it matter whether the factor is thirty or fifty or what?”³ Both China and India are much better off now than in 1949, and the media, not to mention the Pentagon and the State Department, are perennially interested in whether or not the Chinese economy has passed the American economy. We have also made a great deal of progress in data collection and in ways of thinking since Stone wrote, so that we really do have some idea of what the “factors” are. But the uncertainty is still there, especially when we are comparing rich countries with poor countries such as China or India or (even more so) anywhere in Africa. Among the rich countries, whose economic structures are similar to one another, the uncertainty is much smaller and comparisons can be made with some confidence. Market exchange rates are quite close to PPP exchange rates for countries like Canada, the United States, or the nations of Western Europe, and we are on firm ground when making comparisons between them.

Global Growth

Since the end of World War II, which left much of Europe in economic and social disarray, the richer countries of the world have grown rapidly—first repairing the damage and then surging ahead to new levels of prosperity. The rich countries have also grown closer together, and the differences between them are currently small compared with the differences between the group and the rest of the world. Figure 1 shows what has happened to (price-level-adjusted) national incomes for twenty-one rich countries. While measurement is always less than perfect, the data are good and the PPP exchange rates generally reliable for this group of rich countries. The box-and-whisker (or organ pipe) plot is interpreted in the same way as Figure 4 of Chapter 4; the tops and bottoms of the shaded boxes show the




FIGURE 1 GDP per capita in twenty-four rich countries (Australia, Austria, Belgium, Britain, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, Turkey, and United States).

positions of the top and bottom quarters of countries, so that half the countries are in the shaded areas, with the line in the middle indicating the median. The whiskers give an idea of the dispersion of the data, and the dots indicate extreme cases.

The figure shows that other rich countries have shared the slowdown in growth that we have already seen for the United States. The decade of the 1960s was the postwar golden age, with an average growth rate of more than 4 percent a year, a rate that is high enough to increase incomes by a half in ten years. Growth fell to 2.5 percent a year in the 1970s, to 2.2 percent in the 1980s and 1990s, and to less than 1 percent in the decade up to 2010. The extent of the broad decline is exaggerated by, at one end, the catch-up growth after the war—which we would not expect to continue once repairs had been made—and at the other, by the financial crisis. Fixing devastation and destruction, though hard enough, is easier than striking out to levels of income that have not previously been achieved; people remember how things used to be done, and the technology needs to be recreated rather than created from scratch. Once the rebuilding is done, new growth relies on inventing new ways of doing things and putting them into practice, and this turning over of virgin soil is harder than re-plowing an old furrow. Of course, in an interconnected world, innovation can often spread from one country to another—especially to similar countries—so that the burden of invention is spread among many. This interconnectedness, by itself, will tend to speed up growth.

Globalization reduces the costs of moving goods and information from one place to another. It allows goods to be made and, increasingly, services to be performed where it is most efficient and cheapest to do so, and it allows discoveries that are made in one place to be quickly adopted elsewhere. As with new health knowledge or treatments—like knowing the health effects of smoking, or the life-

saving drugs that reduce cholesterol and high blood pressure—discoveries that raised material living standards were quickly internationalized, bringing both health and incomes closer together in the rich countries. For them—where the right political, medical, and economic institutions enabled the changes to be adopted, albeit at different speeds in different countries—we see a remarkable convergence of average incomes, even if the rate of material progress has recently slowed. For these countries, new technology is *reducing* income inequality, just as it reduced health inequality.

The convergence of *average* incomes across these countries tells us nothing about what was happening *within* these countries. Indeed, we have already seen that, in one country, the United States, the growth in average incomes was not widely shared. That countries are drawing closer together does not imply that all of the citizens of the rich world as a whole are drawing closer together. Think of two large crowds that were once apart but are now merging and mixing. If the members of each group are flying apart from one another within their group, the internal divergence can offset and even overcome the merging of the crowds. Seen as a whole, ignoring who is in which country, dispersion might be increasing. I shall come back to this issue when I consider inequality among all the people of the world.

A world of economic growth and narrowing differences between countries is a world to which we have become accustomed, at least if we live in the rich world and if we were born after 1945. High living standards are normal and further growth is confidently expected. Income and health differences between countries have shrunk; travel has become faster, cheaper, and easier; and information is everywhere and instantly available.

The rest of the world doesn't look anything like this. Figure 2 is the same as Figure 1 but includes *all* the countries of the world, both rich and poor. Naturally enough, when we include the poor countries, the




FIGURE 2 GDP per capita in all countries of the world.

range of average incomes is much larger; the boxes are taller, and the whiskers and dots extend farther out. The data are not as reliable, and errors of measurement likely make the spread of incomes appear wider than it really is. More interesting, and less obvious, is that when we look across all countries the spread of average incomes—international country-by-country income inequality—is *not* falling over time. The boxes for 1950 should be ignored for the moment; there are many countries for which there are no data, and many of the omitted countries were very poor, so that the box is too high and too short. After 1950, the difference between the country a quarter of the way up from the bottom and the country a quarter of the way down from the top—the distance between the top and bottom of each box—has remained more or less the same, and if we look at the bottom whiskers we see that the dispersion has actually increased, especially among the poorest countries of the world.

The convergence of average incomes among the rich countries is just what we would expect when growth is driven by new ideas and

new ways of doing things and if new ideas spread rapidly across the world. What is more puzzling is the failure of the poor countries to catch up, which makes Figure 2 look so different from Figure 1. After all, the techniques and knowledge that are the basis for the high living standards in the rich countries are available to the poor countries too. Of course, possession of common knowledge does not imply that all countries should have the same living standards. To be able to use rich-country methods of production requires rich-country infrastructure—roads, railways, telecommunications, factories, and machines—not to mention rich-country educational levels, all of which take time and money to achieve. Yet the gaps between rich and poor provide plenty of incentives to make the investment in that infrastructure and equipment, and, as Robert Solow showed in one of the most famous papers in all of economics, average living standards should draw closer over time.⁴ Why this has not happened is a central question in economics. Perhaps the best answer is that poor countries lack the *institutions*—government capacity, a functioning legal and tax system, security of property rights, and traditions of trust—that are a necessary background for growth to take place.

Growth rates in poor countries have been no lower than, and sometimes higher than, those in the wealthy countries. But while some countries have grown rapidly and are well on the way to catching up, others have fallen further and further behind. The *diversity* of growth experience has been much larger in the poorer countries. Some countries have been able to take advantage of the opportunities for catch-up. A group of Asian countries—China, Hong Kong, Malaysia, Singapore, South Korea, Taiwan, and Thailand—as well as one African country, Botswana, grew at more than 4 percent a year from 1960 to 2010—a more than *sevenfold* increase in average income over five decades. At the same time, the Central Africa Republic, the Democratic Republic of the Congo (DRC), Guinea, Haiti, Mada-

gascar, Nicaragua, and Niger are actually *poorer* in 2010 than they were half a century ago, and there are other countries that almost certainly belong in the group, but for which the data are missing. (Afghanistan, Djibouti, Liberia, Sierra Leone, and Somalia are likely candidates, as are several of the countries that, in 1960, were part of what was then called the Eastern Bloc.) The rapid growth of the successful countries would, by itself, have narrowed the income gaps between countries, but there were enough failures to keep income inequality between countries from falling.

Two of the rapidly growing countries are China and Singapore, but the former has more than three hundred times the population of the latter. India, the other giant, did not pick up speed as early as China, nor has it grown as rapidly, but it has grown at more than twice the world average since 1990. Although China and India are only two countries, their rapid growth at the end of the century meant that around 40 percent of the world's population lived in countries that were growing very rapidly. In contrast, at the “bad” end of the growth distribution, the countries that have been going backwards are in many cases small (though there are exceptions, such as the DRC, which is both large and spectacularly unsuccessful).

Looking at growth rates—not in terms of how many *countries* had high growth, but in terms of how many *people* experienced high growth—global growth takes on a rosier hue. The average *country* grew at 1.5 percent a year in the half-century after 1960, but the average *person* lived in a country that was growing at 3 percent a year. China and India, where so many people live, have grown much faster than the typical country.

One way to think of what has been happening is to imagine a huge crowd of people, the whole population of the world, each carrying the flag of his or her country, like a giant opening ceremony at the Olympic Games. Think of the crowd as marching forward at speeds that are

proportional to the rate at which their incomes are growing—the Indians and Chinese at a run, and some, including the Congolese and Haitians, walking backwards. As we watch the crowd, the two-fifths of all flags that belong to China and India move steadily forward through the crowd, starting near the back (both countries were very poor in 1960) and not yet reaching the front (they are still a long way behind the flags of Europe and North America) but edging closer to the middle. Of course, not everyone in those two countries is moving at the same pace, and, as we shall see, Indians are moving away from other Indians, and Chinese from other Chinese. Yet the rapid average growth of both countries has pulled hundreds of millions out of poverty. And even though all of the countries together are not getting any closer together, the rapid progress of the Indians and Chinese toward the middle has at least made it possible for the entire crowd—the population of the world—to move closer together.

When it comes to a big issue like income inequality among all the people of the world, the word *possible* is a gigantic cop-out. Surely we can do better? The trouble, once again, is the uncertainty in a few key measures. China's rate of growth is one big question; there is a large specialist literature that tries to solve the impenetrable mysteries of Chinese national accounting. Most writers agree that the official rates are too high, but we do not know by exactly how much.⁵ We also don't have very good estimates of China's PPP exchange rate; Chinese (and other) PPPs are subject to a good deal of uncertainty, and the government of China has not participated in all of the price-gathering exercises. If world inequality were expanding or contracting really quickly, measurement uncertainty might not matter. As it is, the truth is that we do not know.

The two largest countries in the world have been among the most successful, at least over the past quarter of a century. Is this *because* they are large, or is it simply that the two countries that did so well

happened to be the largest? Other large countries have also outperformed the world average, at least over some periods, though none have equaled China's staying power. Examples are Brazil, Indonesia, Japan, Russia, and the United States. The "BRIC" countries (Brazil, Russia, India, and China) certainly reap some advantages from their size. A diplomatic corps, a competent bureaucracy, a few well-trained leaders, and the faculty of a world-class university cannot *all* be filled with only a handful of good people, and larger countries have larger pools from which to choose. If scientific discovery—or, more relevant for poorer countries, figuring out how to adapt old knowledge to new conditions—depends on the absolute numbers of scientists or researchers and not on the fraction of the population who are scientists or researchers, then this too gives larger countries an advantage.

In response to a distinguished physicist who once asked me what I was working on, I said the measurement of global poverty. Interesting, he replied: which countries? India, I said, at which point he told me I was talking nonsense: India was one of the most advanced countries of the world. If you count the total number of scientists, and not per capita income or the number of poor people, he was correct, and if scientific work has spillovers that benefit everyone in the country, large countries are blessed. Whether these advantages to scale are enough to elevate growth rates, or whether larger countries tend to grow faster for other reasons, are open questions.

There is much that remains mysterious about why some countries grow rapidly and some grow slowly. In fact, it is not even true that there are perennial fast- and slow-growing countries. At least over the past half-century, the fast-growing countries in one decade have tended not to repeat their performance in the next or subsequent decades. Japan used to be the place that had perpetually high growth, until it didn't any more. India, now one of the most rapidly growing countries, seemed capable of only slow growth for much of its exis-

tence, not to speak of the half-century that preceded its independence, when there was no growth at all.⁶ China is the current long-run superstar, but by historical standards the longevity of its growth spurt is extremely unusual. Economists, international organizations, and other commentators are fond of taking a few high-growth countries and looking for some common feature or policy, which is then held up as the “key to growth”—at least until it fails to open the door to growth somewhere else.⁷ The same goes for attempts to look at countries that have done badly (the “bottom billion”) and divine the causes of their failure.⁸ These attempts are much like trying to figure out the common characteristics of people who bet on the zero just before it came up on a roulette wheel; they do little but disguise our fundamental ignorance.

The foolishness of these schemes reminds me of the search for the keys to success in the Scotland of my youth. Scotland is often cold, wet, and windy, and when I was a child, we knew little about economic growth, and cared less, but we worried all the time about the weather. The long, warm summers of 1955 and 1959 were endless, with golden days spent outdoors in the woods and on the rivers. What were the keys to those successes? I have often thought about it, and for a while had come to the conclusion that those were summers when I was in primary school, and that primary education was perhaps the magic key. But my cousin David, who is a few years older than I am, reminds me that he was in high school in those years, so we gave up on the primary education story. But we remembered that those were both years in which the Conservative party was in power. So perhaps it was not primary education that was the key to success, but politics. All of this is nonsense, but so are all attempts to look at a few successes and a few failures and make fatuous generalizations based on coincidence. Etruscan and Roman haruspices did the same with the entrails of chickens.

Growth, Health, and the Population Explosion

In the six decades since the end of World War II, there have been unprecedented reductions in mortality rates and unprecedented increases in life expectancy (described in Chapter 4), together with the rapid growth in average incomes described earlier. Yet this miracle was far from universally predicted at the time—in fact quite the reverse.

When knowledge of the germ theory of disease brought pest control, clean water, vaccinations, and antibiotics to the poor world, it saved millions of lives, especially the lives of children. Sparing those children brought rapid increases in life expectancy and pulled life chances in poor countries closer to life chances in rich countries. Millions of children lived who would otherwise have died. Alongside the universally welcomed increases in life expectancy came millions of additional people—an increase in global population that was far from universally welcomed. It took most of human history, until the early nineteenth century, for world population to reach a billion people. The two billion mark was reached around 1935, only a century and a quarter later, and three billion by 1960, only thirty-five years later. Instead of slowing down at that point, world population in 1960 was growing at 2.2 percent a year, the highest rate in history and enough to make population double every thirty-two years. The claim that population was exploding was no exaggeration.

The population explosion caused general alarm in the 1960s, among the general public—at least in the rich countries—and also among policymakers, academics, foundations, and international organizations. Much of the concern was humanitarian; many poor countries already appeared to find it hard to feed their populations, so that adding millions more would surely make things even worse. It was as if a poor family, having scraped and scavenged just enough for

a meager meal, found a dozen hungry relatives at the door. Mass starvation loomed. Visitors to India in particular were horrified by the obvious—to them—overpopulation and were appalled at the prospect of adding hundreds of millions more. And indeed, it is easy for first-time Western visitors to be horrified by the poverty and disease in the slums of Delhi or Kolkata; by the beggars, the lepers, the crippled children; by people defecating in the street, and simply by the sheer mass of humanity. How could it be that the addition of yet more people would not make all this even worse?

National security was also a concern. Rising poverty would surely provide a fertile ground for communists—as had already happened in China—and the United States and its allies needed to do what it could to stop the dominoes from falling. Less attractive motives surely also played some role. Concerns about the “quality” of the population had long been a focus of the eugenics movement. Even if eugenic ideas were less popular and respectable after the defeat of Nazi Germany, the possibility that poor, uneducated people should multiply so much more rapidly than rich, educated people seemed to pose a threat to the future of mankind. At the extreme, the enthusiasm for controlling the population of Africa and Asia owed much to the skin color of those whose population was to be brought under control. So it became an object of policy—by countries in their foreign policy, by international organizations in their lending, and by foundations in their giving—to “help” the world’s poor to control their populations by having fewer children. What the world’s poor—the people who were actually having all these babies—thought about all this was not given much consideration.

Why should it have been so universally thought that more people meant poorer people? It may seem obvious that if the world’s food and other goods are shared among more people there will be less for each person. Economists like to call this the lump fallacy: the notion that

there is a fixed “lump” of “stuff,” so that the impoverishment that comes from adding numbers is simply a matter of arithmetic, just as it is for the poor family who have unexpected visitors for dinner. Of course, even this analogy breaks down if the unexpected visitors bring food with them, so that the communal meal may well be better (both nutritionally and socially) than the one that was originally envisaged. The question of impoverishment by numbers is not a question of arithmetic; it is a question of what the new people add, not just what they cost. Perhaps the simplest narrative is that each mouth comes with a pair of hands—too simple for sure, but a better approximation to the truth than the lump story, in which each new person comes with nothing.

We also need to recognize that the African and Asian children who were causing the population explosion were, in the vast majority of cases, children who were wanted by their parents. At the time, even this conclusion was sometimes disputed; people were seen as slaves to their sexual passions, and children were the unfortunate but unavoidable consequence. Certainly, not everyone everywhere has access to convenient modern contraception at low cost, but there is overwhelming evidence—if indeed evidence is required—that on average, if not family by family, people want the children they have, with their own good reasons for having them. The story of unbridled passion provided a convenient rationale for what “we” wanted to do, which was to “help” poor people have fewer children that “they” wanted but “we” did not. No one presented evidence that people wanted such help or that having fewer children would improve their lives. Quite the contrary.

That parents want more children does not necessarily mean that more children are good for society; there may be consequences that parents do not know about, or that they know about but choose to ignore. Children might impose costs on other families. When the

costs are borne by the family itself, we can likely trust the parents to weigh costs and benefits and have children when it makes sense for them to do so. Those children might reduce the resources available to other family members—and indeed, few babies are *born* ready to pull their economic weight. But taking everything into account, including the future economic prospects of the parents and their children as well as the joys of parenthood, we can safely suppose that the additional children add to the family's wellbeing. We might worry that a few parents have children only to exploit or otherwise abuse them, but even this does not make the case that other people will make better choices on their behalf. The more serious argument comes when costs fall on others—more crowded schools or clinics; less common land, firewood, or clean water; or global warming. This argument, often described as the tragedy of the commons, implies that people will have too many children, and it has long been a key plank of the argument for population control.

There are various ways around the tragedy of the commons. Economists like to use prices to solve such problems, and it will sometimes be possible to use a tax to make people pay attention to a social cost that they would otherwise ignore. A classic example is a global tax on carbon, which would do much to combat global warming. But that strategy also illustrates the problem, which is that setting such a tax would require a degree of political agreement that is hard to achieve. Local problems—access to firewood, the use of the commons, who gets water rights—can be dealt with by local political agreement. Although there is never a guarantee that the need for political action will actually be met by the creation of suitable institutions, local political discussions often do solve these disagreements and stop people from doing things that inflict costs on others. The provision of clinics and schools can also be dealt with through local or national politics. The appropriate political institutions may include some sort

of economic or social incentives to limit family size, and this sort of population control—if arrived at in a democratic way—is an appropriate solution to the tragedy of the commons and related difficulties. What such arguments do *not* support is population control by outsiders such as foreign governments, international institutions, or foundations, especially when those organizations have their own interests and an often too limited understanding of the lives of the people they are nominally trying to help.

In the event, much mischief—and worse—was done in the name of international population control. Some of the worst abuses were in India, where voluntary sterilization was often far from voluntary. Although it was Indian politicians and officials who did the abusing, institutions such as the United States Agency for International Development and the World Bank bear much responsibility for encouragement and for intellectual and financial support.⁹ The one-child policy in China, imposed by a nondemocratic government inspired by Western concerns about overpopulation, remains in place today, and it is one of the most serious crimes perpetrated by a modern government against its people. It is not even clear that the policy has been effective: fertility has fallen much further in Taiwan, by exactly the same amount in Thailand, and almost as fast in southern India, and in none of those countries did governments coerce people or violate their deepest and most personal life choices.

In spite of the prophecies of doom, the population explosion failed to plunge the world into famine and destitution. On the contrary, the past half-century has seen not only the decrease in mortality that caused the explosion but also a mass escape from the poverty and deprivation that the population explosion was supposed to cause. What happened, and how could we all have been so wrong?

Not everyone *was* wrong. The economist and demographer Julian Simon consistently challenged the doomsayers, making uncannily

correct predictions about future plenty using arguments that are much more widely accepted today than they were at the time. In his *The Ultimate Resource*, Simon argued that the real source of prosperity is neither land, nor natural resources that might one day be exhausted, but people.¹⁰ Not only does each new mouth bring with it a future worker—which in the long run would make average incomes independent of population size—but it comes with a creative brain. The new ideas that come from some of these new brains are good not only for their owners but also for all of mankind. If having twice as many mouths and twice as many pairs of hands leaves everyone with the same as before, the extra minds can figure out new ways to enable *all* of the hands to do more. Of course, not every newborn is going to be an Einstein, an Edison, or a Henry Ford, nor is every new idea useful for everyone. But because ideas can be shared, not everyone has to be a genius, and any idea that can be applied elsewhere benefits all those who use it, not just the inventor. More children may bring new costs to others—all those schools and clinics—but they also bring benefits in terms of the new ideas and new ways of doing things that are the ultimate basis for economic growth, the tools for the Great Escape. And those benefits might easily outweigh the costs. If so, the world was twice blessed by the health improvements in the 1950s and 1960s—once because of the increase in life expectancy and once because of the explosion of global knowledge and creativity that came with the explosion in numbers.

The economist and demographer David Lam, in his magisterial 2011 presidential address to the Population Association of America, identified the keys to global prosperity in the face of record rates of population growth.¹¹ One key was a fall in fertility: In the face of unprecedented declines in the mortality of their children, families reduced their numbers of children. Parents cared not about the number of children who were *born* but about the number who *survived*.

The children who would have died in previous times no longer “needed” to be born, sparing their mothers the drudgery and dangers of bearing children and sparing their parents the agony of their deaths. We tend to think of the prime beneficiaries of declining child mortality as the millions of children who would otherwise have died but now get the chance to live good lives. True enough, but the parents—particularly mothers—also have their lives transformed; they are free to pursue other activities, such as education and work outside the home. They are also able to devote more resources and time to the nurturing and development of each surviving child.

When fewer children die, parents can reduce the number of children who are born while continuing to have just as many grown-up children to continue the family, to inherit its assets, and to continue its traditions, but at lower costs to themselves in terms of risk and effort. The reduction in fertility did not come immediately—or there would have been no population explosion—but it is visible in the data within a decade or so. As a result, the population explosion was a temporary—albeit a long-lived—event. The world went from a situation in which births and deaths were more or less in balance to one in which births greatly exceeded deaths; and then, after a while, it returned to a balance, but with both births and deaths lower than in 1950. The annual rate of growth of the world’s population, which reached 2.2 percent in 1960, was only half of that in 2011. The people born between the drop in mortality and the drop in fertility made an enormous addition to the world’s population, first as children, when their needs were greater than what they themselves could provide, then as adults, when they were productive and creative, and finally as elderly people, many of whom are retired.

Lam also emphasizes the success of the world economy in responding to the challenge of population growth. This is one of our running themes: that society tends to adapt to new problems, in part by com-

ing up with new ways of doing things—in the case of the population explosion, helped along by all those additional brains—and in part by creating incentives to do things differently. The Green Revolution and other innovations increased the productivity of agriculture and food production grew faster than population. Globalization also helped speed global growth because it allowed production to be carried out in those countries and places where it could be done most efficiently. Limited resources were conserved or substituted for. The price system plays a central role in creating incentives; if nonrenewable resources are becoming uncomfortably scarce, their prices will rise, and people will either make do with less, make substitutions, or direct technical change to find ways of doing without the particular resources altogether.

Economists are often accused of trusting the price system altogether too blindly, and that is sometimes true. Yet both economists and their critics agree that there are great dangers when important resources have no prices, so that, in spite of their value, they are available to users at no cost. Without prices, there is no incentive to economize on such goods. The most important example today is global warming, which, if not addressed somehow, will pose one of the greatest threats to the continuation of the growth of global prosperity.

The misdiagnosis of the population explosion by the vast majority of social scientists and policymakers, and the grave harm that the resultant mistaken policy did to many millions, were among the most serious intellectual and ethical failures of a century in which there were many.

Contraception itself was not the problem. Contraceptives have allowed couples to tailor their fertility to benefit themselves and their children, and the ability to cheaply and effectively regulate fertility has brought better lives to women around the world. As with most innovations, the first beneficiaries were in the rich countries, and this

situation opened up a global inequality. Closing the gap by making the new methods available around the world was rightly seen as a priority and held the potential to do great good in the same way that antibiotics and vaccinations do great good. What *was* deeply wrong was the coercion and the loss of freedom for millions. Rich countries, in the name of helping the poor, were complicit in the coercion. Instead of helping to eliminate a global inequality they added to it, creating greater global injustice. Some harm was done by mistake, and many policymakers and scientists believed that their diagnoses and prescriptions could help poor people. Yet the mistakes were made more likely by the interests of the rich countries themselves—their fear of what life would be like for them in a world with more poor people, and that the population explosion would bolster global communism.

Global Poverty

We have already seen that the growth in national prosperity did little to reduce poverty in the United States, at least after 1975. The world as a whole was more fortunate, and the rapid growth of average incomes, particularly in China and India, and *particularly after 1975*, did much to reduce extreme poverty in the world. In China most of all, but also in India, the escape of hundreds of millions from a traditional and long-established poverty qualifies as the greatest escape of all.

Although the story is clear in broad outline, and its conclusion not under serious challenge, I want to tell it with some care, if only because it is far from obvious how to measure global poverty and what it is we mean when we talk about the number of people living below \$1.00 or \$1.25 a day.

Deciding who is or is not poor is something that is easy for local communities. Development practitioners often conduct “participa-

tory rural assessments,” in which villagers gather around in a central meeting place—under the banyan tree, perhaps—and tell the data collectors about the village, its crops, its main occupations and activities, its water supply, its means of transport, and its people. Those who are seen as poor are often the disabled, or those who are old and have no family to support them. In rich countries too, people are happy to give sensible answers to questions about how much a family needs to “just get by” in their community. Setting a national poverty line is harder, if only because it often comes with an entitlement to special treatment, such as subsidies that are not available to others. Yet, as we saw for the United States, national lines get worked out somehow, and they can be revised or updated in later political debate. Very much the same happened in India, where the original lines, first proposed by academics measuring living standards, were later adopted by the government—in India, the Planning Commission is the guardian of the poverty line—and then revised from time to time—using a favorite device in India, the “expert committee”—whenever the existing lines appeared to be outdated or had stopped attracting widespread support.

The national poverty lines of India and the United States have been produced and discussed within a democratic system, with debate in the press and among interested parties. This gives them the great virtue of domestic legitimacy. Yet many national poverty lines, perhaps even most, cannot be so regarded. For the many governments for whom poverty reduction is of purely rhetorical significance, poverty measurement is done at the instigation of the World Bank, another international agency, or a nongovernmental organization (NGO). Their poverty lines often come, not from domestic debate, but from the guidelines that the World Bank helpfully provides.

Poverty lines constructed by the Bank or using its methods are usually sensible enough, at least as perceived by outside experts.

Indeed, they are most frequently calculated as the amount of income at which a typical family at that income level actually buys a minimally acceptable diet. The fault in these lines is not their lack of *plausibility*, but their lack of *legitimacy*; there is no guarantee that anyone in the country, let alone a poor person in the country, sees the line as a reasonable cut-off point that separates the poor from the non-poor. In effect, these lines are essentially administrative conveniences for the international agencies that need to measure poverty for their own purposes.

The World Bank’s original dollar-a-day global poverty line, and its recent update to \$1.25, came from averaging the national poverty lines of a selection of the poorest countries of the world. These local poverty lines are expressed in local currency, so they have to be converted to a common unit before they can be averaged, and this is done using the PPP exchange rates that I have already discussed. When the Bank first made this calculation twenty years ago, the average was close to a (1985) dollar per person per day, or \$1,460 a year for a family of four; the most recent average—over a different group of countries—is \$1.25 (2005 dollars), or \$1,825 a year for a family of four. At the last stage, this global poverty line, converted back into local currencies, is used in each of the poor countries in the world—the rich countries are excluded from the calculation—to count the number of people living on the local equivalent of the global line. The calculation gives the number of “globally” poor in each country of the world, and these numbers are added up to give totals by region and for the world as a whole.

This calculation has been done on a more or less consistent basis since 1990, and the World Bank now publishes data on global poverty from 1980 to 2008. These are the numbers that I introduced in Chapter 1; Figure 6 of Chapter 1 shows that the number people in the world living on less than a (2005) dollar a day fell from about 1.5 billion

in 1981 to 805 million in 2008. This occurred in spite of an increase in the population in covered countries of nearly two billion people, so that the fraction of the population in poverty fell much faster than the total numbers, from 42 to 14 percent. The decline in numbers is driven almost entirely by the Chinese growth miracle; if China is excluded, 785 million people lived on less than a dollar a day in 1981 compared with 708 million in 2008. This is somewhat less impressive, yet as a fraction of the world's non-Chinese population those that are poor fell from 29 percent to 16 percent.

In India, the other giant modern growth miracle, the number living on less than a dollar a day fell from 296 million to 247 million, and the fraction of the population in poverty from 42 to 21 percent. China and India are the success stories; rapid growth in large countries is an engine that can make a colossal dent in world poverty. The great failure of poverty reduction has been in sub-Saharan Africa. The fraction of the population living on less than a dollar a day in 2008 was 37 percent compared with 43 percent in 1981 and, because Africa's fertility rate has not fallen like Asia's, the number of poor people has almost doubled, from 169 to 303 million.

Africa has a huge land area and looms large on any map, but it is much less densely populated than is either South Asia or East Asia, so that relative failure of poverty reduction in Africa has had less effect on the global poverty count than did the successes in Asia. Even so, we must not make the frequent mistake of somehow discounting the Chinese success. Poverty pessimists, perhaps especially in the aid industry, often make statements to the effect that, with the solitary exception of China, globalization and economic growth have done little to reduce global poverty. But this is exactly the wrong way to think about global poverty. China is not a solitaire, but a country of 1.3 billion people, and to discount their escape from poverty is to argue that Chinese people count for less than Ethiopians, Kenyans,

or Senegalese. Each country is worth study and measurement in its own right, but when we are looking at and trying to measure the well-being of the world, one person has to count the same no matter where he or she lives. There is no premium for living in a small country nor is there a penalty for living in a large one. Global poverty is a cosmopolitan idea and its measurement must be performed on a cosmopolitan basis.

How credible are the poverty numbers? The World Bank measurement scheme that I have outlined is a sensible one—at least apart from its lack of local democratic input—but there are many difficulties along the way; those of us who have been involved in constructing and critiquing these numbers, like many data producers, are much more skeptical and hesitant in their use than most. Even so, I think that we can be confident about the general patterns of global poverty reduction. Rapid growth in China and India is real enough, and the directions of the poverty trends are not affected by possible overstatement of the growth of national income in either country, especially not in China. The African data are often of poor quality, and here there is a good deal more uncertainty, but the stagnation of poverty fits with other things that we know about Africa, such as its relatively slow growth in national income and slow fertility decline. Beyond those broad trends, the rest of the global poverty picture is a good deal murkier.

A weakness of global poverty estimates is their dependence on the PPP exchange rates, so that they are vulnerable to criticism of and uncertainties in those rates—the Marmite problem and more. Another weakness lies in the counts of the numbers of people in each country who are below the line. And there is the question of whether the lines themselves make any sense.

The PPP exchange rates are not calculated every year, but only irregularly; the last three calculations were in 1985, 1993, and 2005,

and the results of a 2011 round are in preparation as I write. Not all countries participate in every year; China, whose size guarantees that it has a big effect on any results, did not participate before 2005, so that the earlier estimates were based on partial information—better than guesses, but far from solidly based. For these reasons, or just because of the difficulty of measurement (we do not know for sure), the global poverty numbers have an alarming tendency to change whenever the PPP rates are revised. These changes have affected the poverty numbers—not just for individual countries, which would be bad enough, but also for whole continents. The 1993 revision suddenly made Africa look much poorer and Latin America much less poor. These were not inconsequential revisions; in sub-Saharan Africa, poverty went from 39 percent to 49 percent.

In 2005, with new data once again, the World Bank upped its estimate of poverty by about a third; of those classified as poor, many more were in Asia than in Africa. Most of this increase came from the Bank changing its poverty line, but the change illustrates the general unreliability of the numbers, not to mention the undesirability of allowing the Bank to be the only source of the numbers on which its own antipoverty efforts are judged. Of course, all of these changes are *statistical* and not real; no one in the world is poorer or richer because the calculations have changed. But those changes can have real effects if international organizations or NGOs shift their efforts (and their rhetoric) to the places where they “see” the highest poverty rates. This is one of many reasons why measurement matters. Much of the recent focus on African poverty postdates the 1993 revision and was arguably influenced by it. Directing aid or attention to the poorest places in the world may be chasing after a chimera, because—to consider a different metaphorical animal—the map of global poverty changes colors like a chameleon.

Trends in global poverty tend not to change very much when the underlying data are revised. Even so, it is possible that the rate of decline of Chinese and Indian poverty may have been understated and that poverty rates are falling faster than the official numbers show. This issue, which is still unsettled, is both technical and deeply political.

It is surprisingly difficult to find out how many people are poor in each country, even once the poverty line is set. The calculations use a *household* survey to ask a random sample of families how much they earn or spend, and then to count the number of people who live in families below the line. There is a cross-check on the surveys from the national income accounts, which provide independent estimates of total spending and total income for the country as a whole. But in many countries, the cross-check fails; the total from the families is often much less than the statisticians think ought to be there, and, worse still, the two totals are drifting further and further apart. Put another way, if we go to households and ask them, they say that their living standards are not improving nearly as rapidly as we would expect from the national growth rate. In one sense, this is similar to what has been happening in the United States; national income is growing, but we see little or no growth for the typical family. The main reason in the United States—that inequality is increasing—is almost certainly *part* of the reason in India and elsewhere. But in India—and to a lesser extent in the United States—the family data are simply *inconsistent* with the aggregate data. This unfortunate gap in the statistics is not unique to India but appears in many countries.¹²

In India, the statistical inconsistency has given rise to a debate that is often vitriolic. On one side—broadly the political right—are those who choose to believe the aggregate data and argue that the poverty measures from the surveys—which are used both by the World Bank

and by the Indian government—understate the decline in poverty. They tell stories of surveyors cheating, sitting under trees or in teashops and making up the data, instead of going to the trouble of asking people. The other—more leftist—side prefers the data from the sample surveys, claiming that if we do not see the poverty reduction when we ask people, then we have no basis for claiming it is there. On their side are the many deficiencies in India's national accounts, as well as the lack of evidence of surveyors sitting in teashops. No doubt both sides have some truth in their claims, yet the debate reminds us that discussions of poverty are sometimes based on facts that are far from solid—fertile ground for people to choose their own preferred version of the truth according to their politics. Underlying all of this is the fact that India's government has become much more pro-business and much less pro-poor in its rhetoric.¹³ Much thus depends on the demonstration that Indian growth is benefiting everyone, and not just the burgeoning and prospering middle class that lives in a few areas of a few cities. Denying the validity of the survey statistics allows those who are doing well not to “see” the poor.

One of my favorite Indian examples illustrates how tiny changes can have huge effects. The great economist and statistician P. C. Mahalanobis, at the Indian Statistical Institute in Kolkata, made many important contributions to the theory and practice of survey design, particularly in household surveys that asked people about their consumer spending. After some experimentation, he decided to ask people how much they had consumed—for example, of rice or wheat—over the past thirty days. In the 1990s, the Indian National Sample Surveys continued to use Mahalanobis's thirty-day rule, although many other countries then used a shorter period of seven days because it was thought that respondents could not remember earlier events with any accuracy. This, it was argued by some, was why so much was being missed, and why poverty was being overstated. The

argument won the day, the switch to a seven-day reporting period was made, and, as expected, average daily expenditures went up. This obscure and technical statistical change *cut the Indian national poverty rate by half*: 175 million people were no longer poor. The length of the reporting period is surely something that only statisticians can get excited about; in current parlance, it is about as deep into the weeds as it is possible to get. Yet a tiny technical issue like this can completely change the measurement and perception of poverty. Poverty reduction is much easier to achieve by statistical means than by actually making people better off!

As a postscript: The Indian change did not stand. After Mahalanobis's experiments were redone, it seemed that the thirty-day period was not terribly inaccurate, and often better than the seven-day period; the thirty-day period was reinstituted in the surveys, much to the delight of the leftists. A more general and more important point is that in India, as in any country where a substantial fraction of the population is poor, there are millions of people who are close to poverty, either just above or just below the line. There are millions who are poor, but who would not be poor if the line were just a little lower, and millions more who are not poor, but who would fall into poverty if the line were just a little higher. As a result, very small changes in the line, or in the way resources are measured, can have very large effects on the number of people classified as poor. Such super-sensitivity undermines the whole business of poverty measurement. We don't really know where the line should be, yet its precise position makes a huge difference. To put it more brutally, the truth is that we have little idea what we are doing, and it is certainly a mistake to let anything important depend on such numbers.

In Charles Dickens's novel *David Copperfield*, the character of Mr. Micawber has his own take on the poverty line. He observes, “Annual income twenty pounds, annual expenditures nineteen pounds

nineteen and six, result happiness. Annual income twenty pounds, annual expenditures twenty pounds ought and six, result misery." One reason that this quote is so memorable is that it is so silly. Why should so much depend on such a tiny difference? And why is someone who is just below the poverty line classified as poor, and worthy of special assistance or the attention of the World Bank, while someone just above the line needs no help and can be left to his or her own devices? When we don't have much idea what the poverty line should be and have great difficulty in measuring income, making such Micawberish judgments is doubly absurd. It makes sense to worry more about people the poorer they are, but not to make sharp distinctions at any critical cut-off.

A final note about the global line. To most people, it is obviously impossible to live in the United States or in Europe on a dollar per person per day. Although no one is expected to do so, and although the United States and Europe are not included in the global counts, this impossibility undermines the validity of the line, even in other countries. After all, millions of people in India live on less than a dollar a day, converted at the PPP exchange rate of about 22 rupees per dollar, and the whole point of these exchange rates is to equalize purchasing power across countries. So if people can live in India on 22 rupees a day—and be far from the worst off—why can't people in the United States live on a dollar a day?

I am not sure that there is a totally convincing answer to this question. The poverty line in India excludes (most of) three things that are important and expensive in the United States: housing, health care, and education. Beyond that, in a warm country like India, there is little need for heating, and much less has to be spent on clothing. People who work near where they live need to spend almost nothing on transportation. If these items are excluded, perhaps an "off-the-

grid" American family of four could buy enough cheap foods—like bulk rice, oatmeal, beans, and a few vegetables—to survive on \$1,460 a year; one recent paper has priced out a "bare-bones" bundle for the United States at around \$1.25 a person a day, or \$1,825 a year for a family of four.¹⁴ Advocates of the validity of the line can also note, correctly, that 22 rupees a day buys a miserable life in India too, and that poor people and their children in India, if not hungry on a daily basis, are among the most malnourished in the world.

Global Income Inequality

It is often claimed that globalization has made the world more unequal, that while the rich have been presented with new opportunities for getting richer, the poor of the world have gained little. These claims have a ring of plausibility. Those of us who are fortunate enough to live in Europe or North America have all of the benefits of the new, interconnected world. At the same time, it is hard to see what good globalization does to the citizens of a poor landlocked country with a poorly educated and unhealthy population.

There are arguments that go the other way. Globalization gives workers in Asia better access to rich-country markets than ever before, and they can do many of the jobs that used to be done in the rich countries, even without being able to migrate. If this happens on a large scale, Asian wages will rise, and American and European wages will fall, narrowing earnings inequality in the world as a whole. The owners of capital also have new investment opportunities with globalization. If capital is relatively plentiful in rich countries and relatively scarce in poor countries, opening up the world allows rich-country capitalists to get richer while poor-country capitalists get poorer. With capitalists getting richer and workers getting poorer,

income inequality will expand in rich countries and contract in poor countries. (Of course, income inequality is not just about the division between workers and capitalists.)

At the beginning of this chapter, I produced data that showed that countries' average incomes are drifting apart, or at the best have shown no tendency to come closer together. Yet some of the world's giant countries are growing very rapidly, so that billions of people now live in countries whose average incomes are closer to middle-class than to impoverishment, and this has been an important force for equalizing incomes across the world. Yet we cannot assess the extent of inequality among all the citizens of the world—what we might call cosmopolitan inequality—using only the averages while ignoring what has been happening to inequality within countries. Just because the Chinese and Indian *averages* are growing very rapidly does not guarantee that the rising tide of prosperity is lifting all boats in China and India. Or, to switch back to my metaphor of the flags at the Olympic Games, the fact that the Chinese and Indian “average” flags are marching from the back of the parade to the middle does not imply that all of the *individual* Chinese and Indian flags are doing the same. The rich Indian tycoons of the high-tech cities might have long reached the front of the parade, leaving poor agricultural laborers right at the back where they have always been. Expanding inequality within countries, if it is severe enough, could offset the giant countries' march to the middle, and cosmopolitan income inequality could be widening.

Chapter 5 documented the recent growth in American inequality. Although the United States is only one country, some of the factors that were important there—new technologies and globalization—must show up elsewhere, or at least in other rich countries. Among the poor countries, there is also evidence that not everyone has benefited from the new opportunities that globalization has brought.

While I do not believe that there is any statement about income inequality that is true in every country of the world—except that it is difficult to measure—it is clear that the general trend has been toward higher income inequality, especially in recent years. The United States is exceptional, both in its level of inequality and in the size of the recent explosion, particularly at the top, but it is certainly not the only country where income inequality is currently increasing. In several of the rich countries, income inequality as measured by the share of the top 1 percent went on falling well into the 1980s, as it had fallen for most of the century, so that the recent uptick was not only smaller than that in the United States but also later.

Chinese economic growth has been geographically uneven, and the cities have done better than the countryside. Such inequalities between rural and urban people create incentives for people to migrate, which would work against huge income differentials, but migration is severely limited in China, and more than a hundred million migrant workers can access good jobs only by being separated from their families. In India, the evidence on expanding inequality is less clear, though again some areas, especially in the south and west of the country, have done better than others. Studies of income-tax records in China and India, part of a multinational research project on top incomes, have found that the share of the top 1 percent has been expanding quite quickly in both countries, though it is only a half (India) or a third (China) the size of the share going to the top 1 percent in the United States.¹⁵ Just to complicate the picture, there is some evidence that inequality is declining in some other large countries, including two traditionally high-inequality countries, Argentina and Brazil.

Many rich countries have also seen increases in income inequality in recent years. Most countries saw a reduction in top incomes in the first half of the twentieth century as wars, inflation, and taxes eroded

the largest fortunes. Over the past few decades, the English-speaking rich countries, like the United States, saw a substantial run-up in the share of the top 1 percent, but this did not happen in the rest of Europe (except for Norway) or in Japan. When the top 1 percent pulls away from everyone else, the bottom 99 percent does less well than the national average. The success of the 1 percent has been different from country to country, which means that the rankings of how well countries have done are sometimes different for the 99 percent than for the whole country.

One interesting comparison is between France and the United States. France has grown somewhat less rapidly than the United States in recent years, but the bottom 99 percent of the population in France saw more rapid growth in average incomes than did the bottom 99 percent of the population in the United States.¹⁶ Or, put the other way around, all but the top 1 percent of the French population did better than all but the top 1 percent of the American population.

The split between English and non-English speakers is what might be expected if top English-speaking managers can sell their services in a world market that is driven by the explosion of top salaries in the United States, but is not open to French, German, or Japanese managers in the same way. A more benign interpretation is that globalization has generated a huge, rich market for top English-speaking managers who, like opera singers or sports stars, now live in an integrated cosmopolitan club of CEOs. In this version of the story, super-sized salaries in the United States and other English-speaking countries are the returns to super-sized talent in the new global market, not the result of American CEOs overpaying themselves and forcing the rest of the English-speaking world to do likewise.

All rich countries face changes in technology as well as competition from lower-wage countries. Not all have shown the same increase in earnings inequality that has occurred in the United States, and

some that seemed originally resistant to the trend are most recently showing an expansion in inequality—a fanning out of earnings—especially above the median. The polarization of jobs and of earnings—with many middle-income jobs being replaced by machines or outsourcing while lower-paid service jobs are doing relatively well—appears to be widespread in rich countries.¹⁷ Polarization, which is new, is limiting the expansion of inequality at the bottom of the income distribution. Other trends—including an increasing prevalence of single-parent families at the bottom and of power couples at the top—are also widespread. Tax and redistribution systems—more comprehensive in Europe than in the United States, and more focused on limiting inequality—seem not to have been able to prevent the recent increase in inequality.

What do these country experiences tell us about inequality in the world as a whole? Are the country expansions in inequality enough to offset the movement of the giants up the world distribution? If the average incomes of countries are moving apart from one another, and if the average country is becoming more unequal, does not that *imply* that the world is becoming more unequal?

Only the last question has a clear answer, which is no. Countries are vastly different in size, and in recent years at least, the giants have grown very rapidly, much more rapidly than average. When we look country by country, we count tiny countries—Guinea-Bissau with one and a half million people—as if they were the same as the giants—India with more than a billion. The fact that Guinea-Bissau and many other small African countries are not doing very well is why *countries* are moving apart, but it tells us nothing about what is happening when we look at whether *people* are moving apart.

What about the contribution of within-country inequality to world inequality? It is important—particularly at the very top of the world income distribution—but not decisive for the vast majority of

people, if only because most inequality in the world comes from differences between countries, not from differences within them. So we are back to the giants—particularly China and India—and how fast they are growing relative to the rest of the world. Growth that is fast enough—even with expanding internal inequality, particularly in China—should sweep everything before it, and the world as a whole should become more equal, at least as long as China remains poorer than the average. Careful estimates, putting all the evidence together, suggest that this is in fact the case and that, in spite of countries pulling apart, and in spite of the growth in internal inequality, global inequality is stable or slowly falling.¹⁸ That may well be correct, though I am not convinced that we know for sure. The big uncertainty is the true growth rates of China and India—whether they are as large as officially stated—and that uncertainty is amplified by the difficulties of making international comparisons between them and other countries.

Finally, we need to ask whether we should care about world inequality and, if so, why. Within a country, inequality tells us something about justice: whether all citizens of a country—who, whether they like it or not, have to pay taxes and conform to the laws and policies of the country—are getting reasonable rewards in line with their obligations. The philosopher Ronald Dworkin has written that “A political community that exercises dominion over its citizens, and demands from them allegiance and obedience to its laws, must take an impartial, objective attitude towards them all.”¹⁹ Admittedly, different people take different views about what justice requires of income distribution, and on whether the high and expanding inequality in the United States is itself unjust, but this is a key part of national discussions about income inequality, whether something ought to be done about it, and if so, just what.

The international situation is different. There is no world government to which people owe allegiance or that has the capability to address international inequalities that might arguably be unjust. The measurement of international inequality is not part of the statistical support for international policy in the way that it is for national policy. Indeed, there are no official statistics on global income inequality among individuals, and perhaps the topic is one that should be left to the curiosity of individual scholars. There is much truth in this, but there are counterarguments too. There may not be a world government, but there are global institutions—for example, the World Trade Organization or the World Bank—whose policies affect the incomes of people in many countries and whose activities are perhaps sufficiently state-like to support justice-based claims by those who are affected. None of these organizations has the authority or the ability to implement a global tax and redistribution system, yet their potential to do good and harm surely makes a case for their at least monitoring income distribution. The world may not be unified but neither is it a set of isolated states that do not interact with one another.