Copyright © 2016 by The President and Fellows of Harvard College All rights reserved Printed in the United States of America

First printing

Library of Congress Cataloging-in-Publication Data

Names: Milanović, Branko, author.
Title: Global inequality : a new approach for the age of globalization / Branko Milanovic.
Description: Cambridge, Massachusetts : The Belknap Press of Harvard University Press, 2016. | Includes bibliographical references and index.
Identifiers: LCCN 2015043601 | ISBN 9780674737136 (alk. paper)
Subjects: LCSH: Equality. | Income distribution. | Globalization—Social aspects. | Globalization—Economic aspects.

Classification: LCC HM821 .M555 2016 | DDC 305—dc23 LC record available at http://lccn.loc.gov/2015043601

Contents

	Acknowledgments vii
	Introduction1
1.	The Rise of the Global Middle Class and Global Plutocrats10
2.	Inequality within Countries
3.	Inequality among Countries
4.	Global Inequality in This Century and the Next155
5.	What Next?

A New Approach for the Age of Globalization

BRANKO MILANOVIC

GLOBAL

INEQUALITY

THE BELKNAP PRESS OF HARVARD UNIVERSITY PRESS

Cambridge, Massachusetts London, England 2016 Introduction

This is a book about global inequality. Throughout the book, I look at both income inequality and political issues related to inequality from a global perspective. Because the world is not united under a single government, however, we cannot dispense with the need to look at individual nation-states. On the contrary, many global issues are played out politically at the level of the nation-state. Thus, greater openness (commercial interchange between individuals from different countries) will have political consequences not at some imaginary worldwide level but within actual countries where the people who are affected by trade live. As a consequence of globalization, for example, Chinese workers might ask for free-trade-union rights from their government, and US workers might ask for protective duties from their government.

Although individual nation-state economies are important, and almost all political action takes place at this level, globalization is an ever stronger force affecting everything from our income levels, our employment prospects, and the extent of our knowledge and information,

1

to the costs of the goods we buy daily and the availability of fresh fruit in the middle of winter. Globalization also introduces new rules of the game through the nascent process of global governance, whether through the World Trade Organization, limits on CO_2 emissions, or crackdowns on international tax evasion.

It is therefore time to look at income inequality not as a national phenomenon only, as has been done for the past century, but as a global one. One reason to do so is simply out of curiosity (a trait much appreciated by Adam Smith)—our abiding interest in how other people, outside our own country, live. But in addition to "mere" curiosity, information about the lives and incomes of others may also serve more pragmatic purposes: it may help us in evaluating what to buy or sell and where, in learning ways to do things better and more efficiently, in making decisions about where to migrate. Or we may use the knowledge acquired from how things are done elsewhere in the world to renegotiate our salary with the boss, to complain about too much cigarette smoke, or to ask the waiter for a doggy bag (a custom that has spread from one country to another).

A second reason to focus on global inequality is that we now have the ability to do so: in the past decade or so, the data required to assess and compare income levels of all individuals in the world have become available for the first time in human history.

But the most important reason, as I believe the reader of this book will appreciate, is that a study of global inequality over the past two centuries, and especially during the past twenty-five years, allows us to see how the world has changed, often in fundamental ways. Shifts in global inequality reflect the economic (and frequently political) rise, stagnation, and decline of countries, changes in inequality levels within countries, and transitions from one social system or political regime to another. The rise of western Europe and North America following the Industrial Revolution has left its imprint on global inequality, driving it up. More recently, the fast growth of several Asian countries has had an equally significant impact, pushing global inequality back down. And national inequality levels, whether increasing in England during the early industrial period or increasing in China and the United States during recent decades, have also had global implications. Reading about global inequality is nothing less than reading about the economic history of the world.

This book opens with the description and analysis of the most significant changes in income distributions that have occurred globally since 1988, using data from household surveys. The year 1988 is a convenient starting point because it coincides almost exactly with the fall of the Berlin Wall and reintegration of the then-communist economies into the world economic system. This event was preceded, just a few years earlier, by a similar reintegration of China. These two political changes are not unrelated to the increased availability of household surveys, which are the key source from which we can glean information about changes in global inequality. Chapter 1 documents in particular (1) the rise of what may be called the "global middle class," most of whom are located in China and other countries in "resurgent Asia," (2) the stagnation of the groups in the rich world that are globally well-off but nationally middle- or lowermiddle class, and (3) the emergence of a global plutocracy. These three salient phenomena of the past quarter century open up several important political questions about the future of democracy, which I address in Chapter 4. But before thinking about the future, we return to the past to understand how global inequality has evolved in the long run.

Global inequality, that is, income inequality among the citizens of the world, can be formally considered as the sum of all national inequalities plus the sum of all gaps in mean incomes among countries. The first component deals with inequality in incomes between rich and poor Americans, rich and poor Mexicans, and so on. The second component deals with income gaps between the United States and Mexico, Spain and Morocco, and so on for all countries in the world. In Chapter 2 we consider within-nation inequalities, and in Chapter 3, among-nation inequalities.

In Chapter 2, I use long-term historical data on income inequality, going back in some cases to the Middle Ages, to reformulate the Kuznets hypothesis, the workhorse of inequality economics. This hypothesis, formulated by Nobel Prize-winning economist Simon Kuznets in the 1950s, states that as countries industrialize and average incomes grow, inequality will at first increase and then decrease, resulting in an inverted-U-shaped curve when one plots inequality level against income. The Kuznets hypothesis has recently been found wanting because of its inability to explain a new phenomenon in the United States and other rich countries: income inequality, which had been decreasing through much of the twentieth century, has recently been on an upswing. This is difficult to reconcile with the Kuznets hypothesis as originally defined: the increase of inequality in the rich world should not have happened.

To explain this recent upswing in inequality, as well as shifts in inequality in the past, going back to the period before the Industrial Revolution, I introduce the concept of Kuznets waves or cycles. Kuznets waves can not only satisfactorily explain the most recent spell of increasing inequality but can also be used to predict inequality's future course in rich countries like the United States or in middle-income countries like China and Brazil. I distinguish between Kuznets cycles as they apply to countries with stagnant incomes (before the Industrial Revolution) and as they apply to countries with steadily rising mean incomes (the modern era). I distinguish between two kinds of forces that drive inequality down: "malign" forces (wars, natural catastrophes, epidemics) and "benign" forces (more widely accessible education, increased social transfers, progressive taxation). I also emphasize the role of wars, which in some instances may be caused by high domestic inequality, insufficient aggregate demand, and search for new sources of profits that require control of other countries. Wars can lead to declines in inequality but also, unfortunately, and more importantly, to declines in mean incomes.

In Chapter 3, the focus is on the differences in mean incomes among countries. Here we face the interesting situation that now, for the first time since the Industrial Revolution two centuries ago, global inequality is not being driven by rising gaps among countries. With the increases of mean incomes in Asian countries, the gaps between countries have actually been narrowing. If this trend of economic convergence continues, not only will it lead to shrinking global inequality but it will, indirectly, also give relatively greater salience to inequalities within nations. In fifty years or so, we might return to the situation that existed in the early nineteenth century, when most of global inequality was due to income differences between rich and poor Britons, rich and poor Russians, or rich and poor Chinese, and not so much to the fact that mean incomes in the West were greater than mean incomes in Asia. Such a world would be very familiar to any reader of Karl Marx, and indeed to any reader of the canonic European literature from the nineteenth century. But we are not there yet. Our world today is still a world in which the place where we were born or where we live matters enormously, determining perhaps as much as two-thirds of our lifetime income. The advantage that people born in wealthier countries possess is what I call "citizenship rent." I discuss at the end of Chapter 3 its significance, its political philosophy implications, and its direct consequence: pressure to migrate from one country to another in search of higher income.

After having looked at the separate components of global inequality, we can return to considering it as a whole. In Chapter 4, I discuss the likely evolution of global inequality in this century and the next. I avoid the seemingly exact projections of global inequality, because in reality they are treacherous: we know that even much more elementary projections of countries' GDPs per capita are most of the time not worth the paper they are written on. It is better, I believe, to try to isolate the key forces (income convergence and Kuznets waves) that are driving nations' and individuals' incomes today and to see where they might lead us in the future. We must remember, though, that in making these predictions, we are often on speculative ground.

While writing Chapter 4, I went back to some of the popular books of the 1970s and 1980s that were trying to predict the future by extrapolating from current trends. I was struck by how time-bound they were, as if imprisoned not only in their space (the place or country where they were written) but even more so in their time.

At the end of À la recherche du temps perdu, Proust marvels at how old people seem to touch, in their own personas, very different epochs through which they have lived. Or as Nirad Chaudhuri writes in the second volume of his beautiful autobiography (Thy Hand, Great Anarch!), it is not impossible to have seen, in one's lifetime, both the peak and the nadir of a civilization-Roman glory at the time of Marcus Aurelius, and the moment when the Forum was abandoned to grass-grazing sheep. Perhaps with age we acquire some wisdom and the ability to compare different epochs that might allow us to better see the future. Yet that wisdom was not evident to me in the writings of the important authors from thirty or forty years ago. It seemed to me that some authors who wrote a century or more ago were more prescient of our dilemmas today than those who were much closer to us in time. Was it because the world dramatically changed in the late 1980s with the rise of China (which nobody writing in the 1970s foresaw) and the end of communism (which similarly was never envisaged)? Can we rule out similarly unexpected events in the next several decades? I do not think so. Yet I hope, though I am far from being certain, that this wisdom of which Proust and Chaudhuri speak and which is acquired with age may be more in evidence to the reader of this book thirty or forty years hence.

I end Chapter 4 with a discussion of three important political dilemmas that face us today: (1) How will China deal with the rising participatory and democratic expectations of its population? (2) How will rich countries manage perhaps several decades of no growth among their middle classes? and (3) Will the rise of the top one-percenters nationally and globally lead to political regimes of plutocracy or, in an attempt to placate the "losers" of globalization, populism?

In the last chapter, I review the main points of the book, distilling its key lessons and making proposals that I believe will be crucial for reducing domestic and global inequalities in this century and the next. For within-national inequalities, I argue for a much greater focus on equalizing endowments (ownership of capital and level of education) rather than on taxation of current income. For global inequality, I argue in favor of faster growth of poorer countries (a rather uncontroversial position) and in favor of lower obstacles to migration (somewhat more controversially). The chapter is divided into ten reflections on globalization and inequality that are more speculative and, unlike the rest of the book, draw more on my opinions than on specific data.

Perhaps the best way to understand the organization of the book and appreciate its symmetry is by means of a schematic chart of its major chapters (Figure I.1).

As the reader can easily see (if she holds a print copy of the book, or if she looks at the total number of words in an electronic copy), this is a relatively short book. It has quite a few graphs, but I hope that they are easy to understand and will help the reader visualize the main points. It is a book that, I believe, can be read with equal appreciation and ease by specialists and by members of the general public, whether well-informed or less-well-informed (even if it is doubtful that anyone would place himself or herself into that last category).



FIGURE I.1. Schematic outline of *Global Inequality*

I owe the reader an explanation about the use of pronouns in the book. I switch quite a lot between the plural *we* and the singular *I*. In general, I use *we* as the usual writer's plural—whenever I think that I am articulating a view that is shared by a significant percentage of economists, social scientists, readers of magazines, or whatever the case may be. Clearly, not everyone whom I embrace under a particular "we" may really hold that opinion. I am aware both of my ascription of opinions to large groups of people and of the fluid nature of the groups themselves. But I try to distinguish this *we* from the *I* that I use when I want to emphasize that some opinions, decisions, ideas, or terms are my own. Thus, to give an example, "we" (that is, economists working on inequality) might think that the Kuznets hypothesis has been discredited by its inability to forecast the recent rise of income inequality in rich countries, but "I" have attempted to redefine it and reformulate it here in such a way that, in the future, "we" may change our opinion about the usefulness of the hypothesis. Yet there is a long way to go before this "I" becomes a "we."

I offer now to the reader the duty—or the pleasure—of taking the first step on the road to the study of global inequality, and perhaps ultimately to global governance, and *the world as one*.

1

The Rise of the Global Middle Class and Global Plutocrats

Intercourse between nations spans the whole globe to such an extent that one may almost say all the world is but a single city in which a permanent fair comprising all commodities is held, so that by means of money all the things produced by the land, animals and human industry can be acquired and enjoyed by any person in his own home.

—geminiano montanari (1683)

Who Has Gained from Globalization?

The gains from globalization are not evenly distributed.

Figure 1.1 shows this phenomenon in a stark way. By plotting percentage gain in income against the original income, we can see which income groups have gained the most in the past few decades. The horizontal axis shows the percentiles of the global income distribution, ranging from the poorest people in the world on the left to the richest (the "global top 1 percent") on the extreme right. (People are ranked by after-tax household per capita income expressed in dollars of equal purchasing power; for details of how income compari-



FIGURE 1.1. Relative gain in real per capita income by global income level, 1988–2008

This graph shows relative (percentage) gain in real household per capita income (measured in 2005 international dollars) between 1988 and 2008 at different points of the global income distribution (ranging from the poorest global ventile, at 5, to the richest global percentile, at 100). Real income gains were greatest among people around the 50th percentile of the global income distribution (the median; at point A) and among the richest (the top 1%; at point C). They were lowest among people who were around the 80th percentile globally (point B), most of whom are in the lower middle class of the rich world. Data source: Lakner and Milanovic (2015).

sons between countries are made, see Excursus 1.1.)¹ The vertical axis shows the cumulative growth in real income (income adjusted for inflation and differences in price levels between the countries) between 1988 and 2008. This twenty-year period coincides almost exactly with the years from the fall of the Berlin Wall to the global financial crisis. It covers the period that may be called "high global-ization," an era that has brought into the ambit of the interdependent world economy first China, with a population of more than one

EXCURSUS 1.1. Where Do the Data for Global Income Distributions Come From?

There is no global household survey of individual incomes in the world. The only way to create a global income distribution is to combine as many national household surveys as possible. Such household surveys select a random sample of households and ask a number of questions on demographics (age, gender, and other characteristics of respondents) and location (where the household lives, including what province, whether in a rural or urban area, and so on), and, for our purposes the most important, questions about the sources and amounts of household income and consumption. Income data include wages, self-employment income, income from ownership of assets (interest, dividends, rental of property), income from production for the household's own consumption (very common in poorer and less monetized economies where households produce their own food), social transfers (government-provided pensions, unemployment benefits), and income deductions such as direct taxes. Consumption data cover money spent on everything from food and housing to entertainment and restaurant services.

Household surveys are the only source of such individualized, detailed information on incomes and expenditures that cover the entire distribution, from the very poor to the very rich. By contrast, data from fiscal sources, such as tax records, generally include only the households of better-off people, that is, those paying income taxes. There are many such households in the United States, but very few in India. Thus, fiscal data cannot be used to generate a worldwide distribution of income.

The size of household surveys varies. Some are large because the country is large: the Indian National Sample Survey includes more than 100,000 households, or more than half a million individuals; the US

Current Population Survey includes more than 200,000 individuals. Many surveys are small, with about 10,000–15,000 people. Such survey data, while never easily available, have recently become more accessible to researchers. For example, in the 1970s and 1980s, not only did relatively few countries conduct surveys, but it was very rare that researchers could get access to "microdata" (that is, individual household data, anonymized to preserve confidentiality). Income distributions were estimated using the government-published fractiles of income recipients (e.g., so many households with incomes between \$x and \$y). More recently, with greater openness of statistical offices and improvements in the processing of large data sets, almost all data, with the notable exception of China, are available at the micro level. This presents significant advantages to researchers: they can redefine income or consumption so as to be comparable across countries or produce inequality measures that are based on households, individuals, or what are called "equivalent units" (adjusting for the fact that larger households enjoy some economies of scale; that is, they do not need a proportional increase in income to be as well-off as smaller households). None of these adjustments is possible without access to the microdata.

The main sources of such microdata are the Luxembourg Income Study (LIS), which includes harmonized survey data (i.e., definitions of income variables that are made as comparable as possible between the countries), mostly from rich countries; the World Bank, which has extensive country coverage and makes some surveys available to outside researchers while other data are available only to World Bank staff; the Social and Economic Database for Latin America and the Caribbean (SEDLAC), located at Universidad de la Plata in Buenos Aires; and the Economic and Research Forum (ERF), located in Cairo, which includes surveys from the Middle East. All of these sources can be easily found on the Internet, but often access to the microdata is restricted to noncommercial uses and "bona fide" researchers, or access is difficult because of the need to know how to download massive databases and apply statistical programs. In addition, for a number of countries (e.g., India, Indonesia, and Thalland), although the data can be accessed directly from statistical offices, that process requires clearance and long waiting periods. So while access to data is becoming much better, it is still not easy. It is also important to realize that even if all the data were suddenly to become easily accessible, factors such as the sheer size of the files, complicated definitions of the variables, and comparability issues mean that income distribution data would never be as simple to use as much more aggregated statistics like Gross National Product.

Now, if each country were to conduct such surveys annually, we could, by collating them, obtain annual estimates of global income distribution. Only rich and middle-income countries have regular annual surveys, however, and even among these countries, annual surveys are something of a novelty. And in many poor countries, especially in Africa, household surveys are done at irregular intervals, on average every three or four years. There are also numerous countries that do surveys only at very long intervals, either because they have no money or technical expertise to field them or because they are at war, civil or foreign. This is the reason why global data can be put together only at approximately five-year intervals (as in this chapter) and are centered around one year, called the "benchmark year," which includes surveys from that year and one or two surrounding years.

National household surveys represent the first building block for determining the global income distribution. The second building block is conversion of such income or consumption data from local currencies into a global currency that should in principle have the same purchasing power everywhere. Why is this important? Because to assess people's incomes and make them comparable, we have to allow for the fact that price levels differ between countries. Thus, to express the real standard of living of people who live in very different environments (countries), not only do we need to convert their incomes into a single currency, but we also have to account for the fact that poorer countries generally have lower price levels. Put in simpler terms, it is less coattain a given standard of living in a poorer than in a richer country: ten dollars will buy more food in India than in Norway. This second building block relies on an exercise called the International Comparison Project (ICP) that is conducted at irregular intervals (the last three rounds were done in 1993, 2005, and 2011) and whose objective is to collect price data in all countries of the world and to use these data to calculate countries' price levels.

The ICP is the single most massive empirical exercise ever conducted in economics. Its final products are the so-called PPP (purchasing power parity) exchange rates. The PPP exchange rate is the exchange rate between, say, the US dollar and the Indian rupee, such that at that exchange rate a person could buy the same amount of goods and services in India as in the United States. To give an example, consider the results for 2011. The market exchange rate was 46 Indian rupees for 1 US dollar. But the estimated PPP exchange rate was 15 rupees per dollar. In other words, if you lived in India, you needed only 15 rupees to buy the same amount of goods and services as a person living in the United States could have bought with 1 dollar. The reason why you needed only 15 rupees (and not 46) is because the price level in India was lower; we can say that it was about one-third (15/46) of the US price level.

It is by applying these PPP exchange rates to the incomes from national household surveys that incomes are converted into PPP (or international) dollars and made comparable across countries. This conversion then enables us to calculate global income distribution. We can see, then, that global income distribution is impossible to calculate without two enormous empirical exercises: hundreds of national household surveys, and individual price data that are aggregated into national price indexes.

However, such massive exercises have their own problems. For household surveys, the most important problem is the imperfect inclusion of people at both ends of the income distribution: the very poor and the very rich. The very poor are omitted because household surveys choose households randomly based on place of residence. Homeless people and institutionalized populations (soldiers, prisoners, and students or workers who live in dormitories) are thus not included, and these people are generally poor. At the other end of the spectrum, the rich tend to underreport their incomes (especially their income from property) and, more alarmingly for researchers analyzing income data, sometimes refuse to participate in surveys altogether. The effect of such refusals on income distribution is difficult to prove directly (because one obviously does not know the income of a household that has refused to be interviewed) but can be estimated from where those who refuse to participate live. It has been estimated that US income inequality might be underestimated by as much as 10 percent because of such nonparticipation (Mistiaen and Ravallion 2006).

These problems are similar or even more serious in other countries and are reflected in two discrepancies between household surveys and macrodata: first, income and consumption reported from household surveys do not fully match household private income and consumption calculated from national accounts (that is, from GDP calculations), and second, statistical discrepancies (called errors and omissions) occur in balance of payments data because of, among other things, money transferred to tax havens (see Zucman 2013, 2015), which, for obvious reasons, is unlikely to be reported in surveys. It is therefore safe to say that household surveys underestimate the number of people who are poor (whatever the definition of poverty) and the number of people who are rich, and their incomes. Lakner and Milanovic (2013) try to adjust globally for the latter, but any such adjustment, while useful, contains a very large degree of arbitrariness due to the simple fact that we know next to nothing about people who refuse to participate in surveys.

The International Comparison Project also suffers from several problems. The most well-known, to which there is no theoretical

solution, is the trade-off between (a) the "sameness" of the baskets of goods and services that are used to measure prices in different countries, and (b) the representativeness of such baskets. To measure differences in price levels, we would ideally like to include the same goods in the "baskets" in all countries. But if we make the baskets exactly the same, we lose representativeness because the staple goods are not the same in all countries. We could achieve identity of baskets by comparing the prices of wine, bread, and beef in all countries, for example, but such a comparison would have little meaning for countries where these items are not widely consumed (e.g., where people consume beer, rice, and fish instead).

It is difficult to find the best solution for this problem, and the ICP at times seems to err in one direction only to then overcompensate by erring in the opposite direction. This produces too much variability in the estimated price levels (see the excellent discussion by Deaton [2005] and Deaton and Aten [2014]). This variability was especially evident for the Asian countries in the last two ICP exercises, in 2005 and 2011. When Chinese or Indian price levels compared to the US price level vary by 20 to 30 percentage points between different rounds of ICP, this produces either much higher or much lower PPP incomes for those countries and thus large swings in the estimates of global inequality. Fortunately for our purposes here, such volatility affects estimated levels of global inequality much more than it affects changes in inequality (up or down) over time.

The data used in this chapter come from more than 600 household surveys covering about 120 countries and more than 90 percent of the world's population over the period 1988–2011. (Most of the data are available on my website: https://www.gc.cuny.edu/Page -Elements /Academics-Research-Centers-Initiatives/Centers-and -Institutes/Luxembourg-Income-Study-Center/Branko-Milanovic, -Senior-Scholar/Datasets.) In the more recent period, after the year 2000, all household survey data are available at the micro level (the level of individual household) with the big exception of China, which does not yet release microdata. All incomes are expressed in 2005 PPP (or international) dollars obtained from the 2005 ICP except where otherwise indicated. Detailed discussion of household surveys and PPPs used is provided by Lakner and Milanovic (2013).

• .

billion people, and then the centrally planned economies of the Soviet Union and Eastern Europe, with about half a billion people. Even India can be included, since, with the reforms in the early 1990s, its economy has become more closely integrated with the rest of the world. This period also saw the communications revolution, which allowed firms to relocate factories to distant countries where they could take advantage of cheap labor without relinquishing control. There was thus a double coincidence of "peripheral" markets opening up and core countries being able to hire labor from these peripheral countries *in situ*. In many respects, the years just before the financial crisis were the most globalized years in human history.²

But the gains, perhaps not unexpectedly in a process of such complexity, were unequally distributed, with some people seeing no gain at all. We focus in Figure 1.1 on three points of interest, where income growth was either the highest or the lowest. They are denoted A, B, and C. Point A is around the median of the global income distribution (the median divides the distribution into two equal parts, each containing 50 percent of the population; one half better-off, the other half worse-off than the people at the median income). People at point A had the highest real income growth: some 80 percent during the twenty-year period. Growth was high, however, not just for those near the median but for a broad swath of people, ranging from those around the 40th global percentile to those around the 60th. This is, of course, one-fifth of the world population.

Who are the people in this group, the obvious beneficiaries of globalization? In nine out of ten cases, they are people from the emerging Asian economies, predominantly China, but also India, Thailand, Vietnam, and Indonesia. They are not the richest people in these countries, because the rich are placed higher in the global income distribution (that is, more to the right in the graph). They are the people around the middle of the distributions in their own countries, and, as we have just seen, in the world, too. Here are some examples of the remarkable cumulative growth experienced by these middle-income groups. The two median deciles (fifth and sixth) in urban China and rural China had their real per capita income multiplied by 3 and about 2.2, respectively, between 1988 and 2008. For Indonesia, median urban incomes almost doubled, and rural incomes increased by 80 percent.³ In Vietnam and Thailand (where the population is not split into rural and urban), real incomes around the medians more than doubled.⁴ These groups were the main "winners" of globalization between 1988 and 2008. For convenience, we call them the "emerging global middle class"-although, as I shall explain later, because they are still relatively poor compared with the Western middle classes, one should not assign to the term the same middle-class status (in terms of income and education) that we tend to associate with the middle classes in rich countries.

Let us move now to point B. The first thing to notice is that it is to the right of point A, meaning that people at point B are richer than people at point A. But we also notice that the value on the vertical axis at point B is nearly zero, indicating the absence of any growth in real income over twenty years. Who are the people in this group? They are almost all from the rich economies of the OECD (Organization for Economic Cooperation and Development).⁵ If we disregard those among them who are from the relatively recent OECD members (several Eastern European countries, Chile, and Mexico), about three-quarters of the people in this group are citizens of the

"old-rich" countries of Western Europe, North America, Oceania (the three areas are sometimes represented by the acronym WENAO), and Japan. In the same way that China dominates at point A, so do the United States, Japan, and Germany dominate at point B. People at point B generally belong to the lower halves of their countries' income distributions. They are from the bottom five deciles in Germany, which from 1988 to 2008 managed cumulative growth of only between 0 and 7 percent; from the lower half of the US income distribution, which experienced real growth of between 21 and 23 percent; and from the lower deciles in Japan, which saw either a decline of real income or overall growth of 3 to 4 percent. For simplicity, these people may be called the "lower middle class of the rich world." And they are certainly not the winners of globalization.

It is simply by contrasting the groups at these two points that we have established empirically something that has been felt by many people and widely discussed in economic literature as well as in public fora. We have also highlighted one of the key issues of the current globalization process: the diverging economic trajectories of people in the old rich world versus those in resurgent Asia. In short: the great winners have been the Asian poor and middle classes; the great losers, the lower middle classes of the rich world.

Such a bald statement may not surprise many people today, but it would certainly have been surprising to many if it had been made in the late 1980s. Politicians in the West who pushed for greater reliance on markets in their own economies and the world after the Reagan-Thatcher revolution could hardly have expected that the muchvaunted globalization would fail to deliver palpable benefits to the majority of their citizens-that is, precisely to those whom they were trying to convince of the advantages of neoliberal policies compared with more protectionist welfare regimes.

But such a statement would appear even more surprising to those, including the Nobel Prize-winning economist Gunnar Myrdal, who

worried in the late 1960s that the Asian masses, numbering many millions and barely able to survive on their low incomes, would remain mired in perpetual poverty. An entire literature of the 1950s and 1960s (such as Paul Ehrlich's The Population Bomb [1968]) had as its main theme the dangers that population growth presented for economic development in the Third World. The Asian experience of the last quarter of the twentieth century has fully contradicted such dire warnings. Instead of the "Asian Drama," which was the title of Myrdal's book, we hear today about the East Asian Miracle, the Chinese Dream, and Shining India, all coined to parallel the American Dream and the German Wirtschaftswunder (economic miracle).

I point to this example here, very early in the book, to highlight the difficulties that beset any long-run forecasting of economic development, particularly on a global scale. The number of variables that can and do change, the role of people in history ("free will"), and the influence of wars and natural catastrophes are so great that even forecasts of broad tendencies made by the best minds of a generation are seldom correct. We should be aware of that difficulty when in Chapter 4 we discuss the likely economic and political evolution of the world in the rest of this century and the next.

The contrast between the fortunes of the two middle classes illustrates one of the key political questions today: are the gains of the middle class in Asia related to the losses of the lower middle class of the rich world? Or, to put it differently, is the stagnation of incomes (and wages, since wages account for the lion's share of income of the lower middle and the middle class) in the West a result of the success of the Asian middle class? If this wave of globalization is holding back the income growth of the rich world's middle classes, what will be the result of the next wave, involving ever-poorer and more populous countries such as Bangladesh, Burma, and Ethiopia?

Let us now go back to Figure 1.1 and look at point C. Its interpretation is simple: we are dealing here with the people who are globally

21

very rich (the global top 1 percent) and whose real incomes have risen substantially between 1988 and 2008. They too are the winners of globalization, almost as much as (and as we shall see in a moment, in absolute terms even more than) the Asian middle classes. People who belong to the global top 1 percent are overwhelmingly from the rich economies. The United States dominates there: half of the people in the global top 1 percent are American. (This means that approximately 12 percent of Americans are part of the global top 1 percent.)⁶ The rest are almost entirely from Western Europe, Japan, and Oceania. Of the remainder, Brazil, South Africa, and Russia each contribute 1 percent of their populations. We can call those in group C the "global plutocrats."

Comparison of groups B and C allows us to address another important cleavage. We have seen that group B, with zero or negligible gains from globalization, consists mainly of the lower middle class and the poorer segments of the rich countries' populations. In contrast, group C, the winners of globalization, consists of the richer classes from these same countries. An obvious implication is that the income gaps between the top and bottom have widened in the rich world, and that globalization has favored those in the rich countries who were already better-off. This too is not entirely surprising, since it is generally acknowledged that within-nation inequalities in the rich world have increased during the past twenty-five to thirty years.⁷ This is the topic we shall address in Chapter 2. But what is important, and rewarding in an epistemological sense, is to see that these effects are observable when we look at the world as a whole, too.

Figure 1.1 displays only a very rough image of the winners and losers of globalization. Many additional ways to look at these data are possible: we could look in much more minute detail at the horizontal axis (splicing the world's population into smaller "fractiles" of, say, 1 percent), or we could look at how given income groups (such as the poorest 10 percent of people in China versus the poorest 10 percent of people in Argentina) have fared over the same twenty years, or we could define income gains in standard exchange-rate dollars rather than adjusting them to take into account different price levels in different countries. But whatever adjustment we make, the essential shape of the gains and losses shown here does not change: it always appears as a reclining S curve (or what some people have called an "elephant curve," because it resembles an elephant with a raised trunk). The percentage gains are always the strongest among the middle classes in emerging economies and the global 1 percent; they are always the least among people situated around the 75–90th percentile of the global income distribution, in other words, the middle and lower middle classes in OECD countries.⁸

This shape, with a trough at the position of the relatively well-off percentiles, is very unusual in the case of individual countries. Normally, graphs such as these, which are called growth incidence curves (GIC), either rise more or less continuously, indicating that the rich have gained more than the poor, or, on the contrary, slope downward continuously, demonstrating the reverse. A reclining S curve shows that the changes in income have been such that the rich and the middle class have benefited more than those in between. Within an individual country, such changes are not likely because they would imply that either economic policies or technological change had been "calibrated" in such a way as to benefit the top 1 percent or 5 percent, to go against the interests of those placed immediately below, and then to benefit those further down. Such discontinuities are not very likely to occur in the way either new technologies or new economic policies help or hamper various income groups. For example, it is not probable that a policy that cut marginal tax rates for the top 5 percent would be accompanied by another policy that increased taxes on those just below the top 5 percent level. Here, however, we are dealing not with a single country distribution but with a global distribution that is the product of several factors: (a) the differences in countries'

growth rates (or to be more specific: China's faster growth rate in comparison with that of the United States), (b) countries' original positions in the global income distribution in 1988 (when China was so much poorer than the United States), and finally (c) changes in the countries' own income distributions, which are affected not only by domestic policies but by globalization (principally by China exporting cheap goods to the United States). These factors explain how such unusually shaped curves, like the reclining S curve, are possible. What do we expect the shape of the global incidence curve to look like in the next thirty years? We shall address this issue in Chapter 4.

A very important caveat regarding the interpretation of "winners" and "losers" and of the meaning of the reclining S curve is that so far we have dealt only with relative gains across the global income distribution. The vertical axis in Figure 1.1 shows the cumulative percentage change in real income between 1988 and 2008. How would the results look if instead of *relative* change (percentage gain) we considered *absolute* change (number of dollars gained)? As we shall see, this change in perspective alters the results in a rather dramatic way.

Absolute Income Gains along the Global Income Distribution

Suppose that we take the entire increment in global income between 1988 and 2008 and call it 100: Figure 1.2 shows that 44 percent of the absolute gain has gone into the hands of the richest 5 percent of people globally, with almost one-fifth of the total increment received by the top 1 percent.⁹ In contrast, people whom we have termed the main beneficiaries of the current era of globalization, the "emerging global middle class" have only received (by ventile) between 2 and 4 percent of the increase in the global pie, or in total about 12–13 percent.



FIGURE 1.2. Percentage of absolute gain in real per capita income received, by global income level, 1988–2008

This graph shows the percentage of total absolute gain in real household per capita income (measured in 2005 international dollars) between 1988 and 2008 received by groups at different points of the global income distribution. We take the increase in total world real income as 100 and calculate how much of it was received by different ventiles (groups of 5% of the population) or percentiles of the global income distribution. The graph shows that the absolute gains in income went mostly to the richest 5% of the world population. The top 1% got 19% of the total global income increase. Data source: Lakner and Milanovic (2015).

How is this possible, and does this distribution of absolute gains invalidate our previous point regarding the winners and losers? It is possible simply because of the enormous gaps in real income that exist between the top, the median, and the bottom of the global income distribution. In 2008, the average per capita disposable (after-tax) income of the global top 1 percent was just over \$71,000 per year, income at the median was around \$1,400, and people who were in the poorest global decile had annual incomes under \$450 (all figures are in 2005 international dollars). In looking at these numbers, we immediately see that what is but a rounding error for the incomes at the top is equivalent to the entire annual income of the poor! Now, it is clear that a very small percentage gain at the top, or around the top, will represent a huge share of the overall absolute gain. Suppose, for example that the income of the richest 1 percent increases by only 1 percent, or \$710. But that amount represents onehalf of the total income of the people at the global median. This is why both the large relative gains at the very top (the income of the top 1 percent grew by two-thirds between 1988 and 2008) and the almost nonexistent gains among the lower middle classes of the rich world (whose incomes increased by only 1 percent), when translated into absolute gains, look so remarkable compared with the absolute gains of the emerging global middle class. It is just a very good illustration of how hugely unequal is the distribution of incomes globally.

Does this skewed distribution in absolute gains make us revise our previous conclusion regarding the winners and losers? It does not. Rather, in some respects it emphasizes what we concluded for the richest 1 percent or 5 percent, because their considerable percentage gains appear even more stunning when we look at them in absolute amounts. (For more on absolute versus relative measures, see Excursis 1.2.) It does not make us revise our conclusion for the lower middle classes of the rich world, either, because they, like most of us, look primarily to their percentage gains (which were minimal), and when they compare their position with that of others, they are likely to contrast it with the real percentage gains realized by the top. So their income stagnation is very real. And, finally, it does not affect our conclusion about the success of the Asian middle classes either, because they too are likely to consider their relative gains first. But the introduction of the absolute measurement allows us to look at the same data from a different angle and to better perceive the immense differences in income that exist in the world today. It also highlights an important point: we should not conflate the middle classes from

EXCURSUS 1.2. Absolute versus Relative Measures of Income Inequality

In addition to highlighting the massive income gaps in the world, the comparison of relative and absolute gains in income has another value relating to the decades-old discussion of relative versus absolute measures in income distribution studies. Almost all of our inequality measures are relative, in the sense that if everybody's income increases by the same percentage, inequality is deemed unchanged. But an equal percentage increase for all corresponds to absolute gains that may be extremely unequal: a person who started the race with an income one hundred times higher will also have absolute gains that are one hundred times greater. So why are relative measures better?

First, relative income measures are conservative because they show no change in inequality in cases where absolute measures would show an increase (when all incomes go up by the same percentage) or a decrease (when they all go down by the same percentage). On inequality, which is a topic of considerable moral and political importance, and at times a very inflammatory topic indeed, we do not want to err in the direction of inflaming it further. Conservatism (in terms of measurement, not necessarily in terms of policy) is to be preferred.

Second, one of the disadvantages of absolute measures is that they are bound to increase with practically any increase in the mean: when incomes rise, the absolute distance between the rich, the middle class, and the poor becomes greater even if the relative gaps remain the same. Think of the distribution as a balloon. As the balloon expands, the absolute distance between the points on the balloon increases. Focus on absolute distances presents the disadvantage that practically every increase in the mean (blowing up the balloon) could be judged to be pro-inequality. We would lose the sharpness with which we can currently distinguish between pro-poor and pro-rich growth episodes. With an absolute inequality criterion, it would be hard to argue that the United States entered a period of rising inequality after the 1980s (a topic which we address in Chapter 2). Since growth in the 1960s was strong, it is very likely that the absolute gaps increased then, too. So would we say that inequality in the United States started rising in 1945, or even earlier, and has not stopped since? But clearly these different periods were not the same as far as inequality is concerned.

Third, inequality and income growth are just two manifestations of the same phenomenon. Again, this point is most obvious in global inequality studies, where changes in total inequality among world citizens depend crucially on the growth rates of different countries. For the more mathematically minded, it may be easier to see this fundamental similarity between inequality and growth by thinking of the mean income as the first moment of a distribution, and of inequality as the second moment of a distribution (the variance). Growth is simply the relative increase in the first moment, and inequality is the relative increase in the second moment. The measures that we use to assess success or failure in economic development (relative change in GDP per capita) should be related to the measures we use to assess success or failure in distribution of resources (relative change in a measure of inequality). Focus on the absolutes in growth, as in inequality, would lead us to nearly always find that growth in rich countries, however small in percentage terms, would be greater than growth in poor countries, however huge. If the United States grew by 0.1 percent per capita annually, that growth would increase the absolute GDP per capita of each American by about \$500, which is more than the GDP per capita of many African nations. Should we then deem Congo, in any given year, to have been as successful as the United States only if it doubles its per capita income---a feat that no human community has ever achieved in recorded history? So the logic of relativity that applies to growth should also apply to inequality.

A final argument is that relative increase in income correlates with gains in utility if we believe that personal utility functions are logarithmic in income—that for a person whose income is \$10,000 to experience the same increase in welfare as a person whose income is \$1,000, the absolute income gain ought to be ten times greater. In other words, one additional dollar will yield less utility, or seem less important, to a rich person than to a poor person. If we think that this is a reasonable assumption, we can then also interpret the data given in the growth incidence curve as changes in utility: an 80 percent income increase around the global median adds to the utility of people there more than a 5 to 10 percent increase in real income adds to the utility of the lower middle classes in rich countries (even if the absolute dollar gains of the latter may be larger). By this route too, we come to the conclusion that relative income changes are a more reasonable metric than absolute income changes.

the emerging market economies (people with per capita incomes of approximately between \$1,000 and less than \$2,000 per year) with the lower middle classes of the rich world (people with after tax incomes of approximately \$5,000 to \$10,000 per year; all in 2005 international dollars).

Comparison of Figure 1.1 (relative income gain) and Figure 1.2 (absolute income gain) highlights a feature that we shall often find when we analyze the changes brought about by globalization: we will very seldom be able to point to a change that has either wholly positive or wholly negative effects, or that is entirely unambiguous in its effects on all people, or in all its manifestations. In this case, we see that the much greater relative income gains for the middle classes of the emerging market economies did not always translate into greater absolute gains. By their very nature, dramatic economic movements affect various countries and groups of people differently, so that even in the case of a change that we might view as overwhelmingly positive, certain people and groups would be made worse off by it.

It is this fundamentally ambivalent nature of globalization that I hope to bring out in this book. The reader needs to be constantly aware that globalization is a force both for good and bad. Ideally, he or she, even when reading about some aspects that seem "good," should be on alert for thinking about what drawbacks or "bad" effects may lurk behind them (and conversely, when reading about "bad" effects). Our ability to comprehend and include all the "goods" and all the "bads" and to give them a subjective weighting will, in the last analysis, determine how we feel about globalization. But it is precisely this ambivalence, combined with the fact that our personal weighting schemes are by necessity different—not only because we might believe in different things, but because we ourselves or people we care about may be affected positively or negatively by globalization—that will make unanimity about the effects of globalization forever elusive.

The Effects of the Financial Crisis

We have so far discussed the changes between 1998 and 2008 because they best represent the effects of "high globalization" and because our data for that period have been well organized and made as comparable as possible. But new data and information from 2008 to 2011 are now available. In most respects, this last short period—which comes just after the financial crisis—is a continuation and even an acceleration of the globalization trends described above; but it continues the trends with a twist.

A trend that became even stronger in 2008–2011 was the growth of the global middle class, fueled during these three years, as in the previous twenty, by high growth rates in China. Between 2008 and 2011, the average urban income in China doubled, and rural incomes increased by 80 percent, driving the global growth incidence curve around the median substantially above its 1988–2008 point. Thus the growth of the global middle class became even more visible and entrenched (see Figure 1.3).

On the other hand, the absence of growth in the rich world meant not only that the incomes of the lower middle classes in these countries continued to stagnate but also that the stagnation extended toward the top. There, too, there was no growth, and this is why point C has remained where it was in 2008 (compare Figures 1.1 and 1.3).¹⁰

The effect of the financial crisis on the global distribution of incomes is not surprising. What is unclear is how significant a break in global economic history this crisis, often referred to as a global financial crisis, represents. First, it should be noted that the very term



FIGURE 1.3. Relative gain in real per capita income by global income level, 1988–2008 and 1988–2011

This graph shows relative (percentage) gain in real household per capita income (measured in 2011 international dollars) at different points of the global income distribution for two different time periods: 1988–2008 (replicating the graph in Figure 1.1, except that we now use 2011 instead of 2005 international dollars) and 1988–2011. We see the continuation of very strong gains around the middle of the global income distribution but a slowdown of gains among the global top 1%. Data sources: Lakner and Milanovic (2015) and author's data.

"global" is a misnomer because the slowdown (or the recession) affected, at first, only the rich economies. It should more properly be labeled a recession among the Atlantic economies. Second, the longterm evolution of incomes at the level of nations, that is, the rebalancing of economic activity in favor of Asia and away from Europe and North America, was not interrupted but rather was reinforced by the crisis. Thus, the crisis represented not a break in this trend, but rather the reverse: reinforcement of an already existing trend. Third, the rebalancing has a counterpart in the distribution of personal incomes worldwide in the sense that it changed the shape of the global income distribution from being strongly twin-peaked (having many people at very low incomes, then practically nobody in the middle, and finally more people at very high income levels) to being fuller in the middle, such that the global income distribution is now beginning to look like the distribution of a single country. We are, of course, still far from that point, but we are certainly closer to it in 2011 (or today) than we were in 1988. This trend, too, was merely reinforced during the crisis.

Figure 1.4, which shows the distribution of world population according to income level in 1988 and 2011, illustrates very clearly the emergence of the global middle class and the diminution (flattening) of the two-humped shape of the global income distribution. What is interesting, however, is that an "emptiness in the middle" still largely characterizes the distribution of world population according to the mean income (or GDP per capita) of the country where people live, as can be seen in Figure 1.5. The contrast between the two figures illuminates the fact that while India and Indonesia, and to a somewhat lesser extent, China, remain poor countries judged by their mean incomes, income distributions in these countries are sufficiently wide and skewed to the right that a significant number of their citizens are now filling that space, the empty middle that used to exist between the two peaks.





This graph shows the distribution of world population according to real household per capita income (measured in international dollars) in 1988 and 2011, based on household surveys. The area beneath each curve is equal to total world population, respectively, in 1988 and 2011. Between 1988 and 2011, there was an expansion in the proportion of people with incomes around the middle (the "global middle class"). The graph shows that this global middle class is still relatively poor by Western standards. Data sources: Lakner and Milanovic (2015) and author's data.

The evolution of incomes in China is here again emblematic of global changes, perhaps because the increase was the fastest of any country and involved the most people. According to the household survey data for 2011, mean income in urban China has, for the first time, caught up with and even exceeded mean incomes in several European Union (EU) member countries. Urban China now has a higher mean income (in PPP terms) than Romania, Latvia, or Lithuania. In 2013, China's GDP per capita was still lower than that of the poorest EU members (Romania and Bulgaria), but the gap was less than 30 percent, and with the currently expected rates of growth, by



FIGURE 1.5. Distribution of world population by real GDP per capita of the country in which people live (year 2013)

This graph shows how world population would be distributed if we assigned to people their countries' mean income (GDP per capita) instead of their actual per capita income (as in Figure 1.4). Labels show selected countries. We see that there are relatively few people living in countries with "middling" levels of income. Data source: Calculated from the World Bank's World Development Indicators (WDI) database (http://data.worldbank .org/data-catalog/world-development-indicators, version September 2014).

the time the reader holds this book in his or her hands, China's GDP per capita will undoubtedly have reached the level of the poorest EU countries.¹¹ This is an epochal change, for although Romania, Bulgaria, and the Balkans have been the poorest part of Europe since the Middle Ages, their per capita incomes in the late nineteenth century were twice as high as China's.¹² Moreover, since we can expect that China will continue to grow faster than the core EU countries, even if its growth rate decelerates, its mean income will catch up with the EU average in another three decades.¹³ This would be, in a historically very short period, a remarkable reversal of fortunes, or rather a

return to a pattern of distribution characteristic of economic activity in the Eurasian space several centuries ago: per capita incomes may once again be highest in two coastal regions, one facing the Atlantic (western Europe) and the other facing the Pacific (China), while they are lowest in the hinterland of Eurasia. Peninsular Europe's exceptionalism will have come to an end.¹⁴

Another way to look at the change in incomes over the past several decades is to compare the mean income of people in the lower part of the US income distribution with that of people who are relatively well off in urban China (Figure 1.6). Note that since practically all of the United States is urbanized, we are de facto comparing urban



FIGURE 1.6. The convergence of Chinese and US incomes, 1988-2011

This graph shows the change in annual real household per capita after-tax income (measured in 2005 international dollars) between 1988 and 2011 for people in the US second decile and the Chinese eighth urban decile (based on household survey data). Vertical axis is in logs. Although the US second decile (while relatively poor by US standards) was still better-off than the Chinese eighth urban decile in 2011, the gap between the two has been diminishing. Data source: Author's data. United States with urban China. The catch-up between 1988 and 2011 is quite apparent. The gap in real incomes decreased from more than 6.5 to 1 to only 1.3 to 1. (This catch-up could be illustrated by using other parts of American and Chinese distributions, but it is more striking in this example because the two income levels are becoming similar. If we used higher parts of the US distribution, the gaps would still have been very large.) There is also no doubt that this diminution of the gaps in per capita household incomes corresponds to a diminution in the real wages gap.

The Global Top 1 Percent

We have seen that although the global top 1 percent had a very good run between 1988 and 2008, their fortunes darkened between 2008 and 2011. The reason is simple: most of the people in the global top 1 percent belong to the high parts of income distributions in the rich countries (for example, 12 percent of the richest Americans are in the global top 1 percent), and their income growth slowed down or was brought to a halt by the financial crisis. This slowdown might seem surprising at first sight, given the tremendous increase in interest, awareness, and concern with top incomes in the rich world, and especially in the United States. But the contrast between the huge interest in top incomes and simultaneous slowdown in their growth is explained in part by the fact that while most incomes in rich countries declined during the crisis, top incomes remained stable or declined less. Although remaining stable might appear "good" (or perhaps even "unfair" from the point of view of other people in rich countries), it was not good enough for the global top 1 percent to maintain as high a position in comparison to the global median as before the crisis. This is because the median and the mean global income have continued to grow.

Another reason for the contrast between the recent slow growth among the global top 1 percent and popular concern with inequality is that the growth on the top was much more concentrated among the super-rich than before. In effect, if we want to focus on those who continued to gain throughout the crisis we should focus not on the global top 1 percent (which includes some 70 million people, about equal to the population of France) but on a much narrower group of super-wealthy individuals. There are, of course, many fewer of these individuals, and they are not included in household surveys.¹⁵ We shall look at them very briefly in the next section, using an entirely different data source, *Forbes*'s list of billionaires. The list includes in 2013 and 2014 about 1,500 individuals who together with their families represent one-hundredth of one percent of the world population (yes, it is 1 percent of 1 percent).

Let us first return to the global top 1 percent as represented in household surveys. Figure 1.7 shows the countries that have more than 1 percent of their population in the global top 1 percent. We have already seen that the United States is very well represented, with 12 percent of its population being in the global top 1 percent and accounting for about half of all the people there. Other large advanced economies, like Japan, France, and the United Kingdom, have between 3 and 7 percent of their populations in the global top 1 percent, while Germany has only 2 percent. Not shown in the graph are Brazil, Russia, and South Africa, whose top one-percenters are also in the global top 1 percent. But this is not the case for China and India, who have fewer than 1 percent of their populations in the global top 1 percent. The global top 1 percent is thus heavily dominated by the old-rich countries: China's upward march through the global income distribution has not yet spread, in sufficient numbers, to the very top.16

The income share of the global top 1 percent in 2008 was 15.7 percent. This number represents their share of *global* disposable income. It can be compared with *national* top 1 percent shares reported in the World Top Incomes Database (WTID), but one has to be aware that the incomes reported in WTID are before transfers and taxes and across





This graph shows the countries that have more than 1% of their population in the global top 1%. We see that 12% of the richest Americans belong to the global top 1%. Country abbreviations: CAN Canada, CHE Switzerland, CYP Cypress, DEU Germany, FRA France, GBR United Kingdom, IRL Ireland, JPN Japan, KOR South Korea, LUX Luxembourg, NLD Netherlands, NOR Norway, SGP Singapore, TWN Taiwan, USA United States. Data source: Lakner and Milanovic (2013).

fiscal units, while incomes discussed here are after taxes and are calculated across individuals.¹⁷ (Fiscal data cannot be used to calculate the top 1 percent share globally because fiscal data are available only for a relatively small subset of countries.) The biggest difference between the two data sources is the use of market, that is, pre-transfer and pre-tax, income by WTID rather than disposable, that is after-tax, income as used in household surveys. The share of the top 1 percent will always be greater in terms of market income than disposable income because government redistribution reduces inequality. For example, redistribution via government transfers and direct taxes in the United States in 2010 reduced the share of the top 1 percent from 9.4 percent of total market (or "pre-fisc") income to less than 7 percent of total disposable income.¹⁸ (It should also be mentioned

that the people who are in the top 1 percent according to pre-fisc income are not necessarily the same people who are in the top 1 percent according to disposable, that is post-fisc, income.) Using the United States as a comparator, we can say that the share of the global top 1 percent in world income is more than twice as high as the share of the top 1 percent in US total income (15.7 versus less than 7). This gives us a fair shorthand view of how high the concentration of income is at the global level. Yet another, more focused, view is provided by the *Forbes* annual list of billionaires.

Note, however, that when we discuss Forbes's list of billionaires, we are making an important methodological move: instead of looking, as we have done so far, at incomes or consumption, which are annual flow variables, we are looking at wealth, which is a stock variable (that is, measured at one point in time) and is the result of accumulation of savings, returns on investment, and inheritance over the years. Wealth inequality is greater than inequality of income or consumption in almost every country. Not only are there tiny groups of enormously wealthy people-a phenomenon on which we shall focus in the next section-but even in the advanced countries (say, the United States or Germany), between a quarter and one-third of the population has negative or zero net wealth.¹⁹ But very few people in these countries have zero income, and no one has zero consumption. Thus it can be seen even at an intuitive level that wealth must be much more unequally distributed than income or consumption, and that comparisons between wealth inequality and income inequality have to be made very carefully.²⁰ It is because the wealth data for the super-rich are of better quality (and to some extent more revealing) than the income data for the top 1 percent that we use wealth data rather than income or consumption data to shed light on the position of the super-wealthy.²¹

To see the difference between income and wealth distributions on the global level, consider Table 1.1, which shows estimates of the in-

come and wealth shares of the global top 1 percent. For income, we have three estimates: first, the conservative one, based on household surveys alone, which (as discussed in Excursus 1.1) tend to miss the richest people and thus underestimate the share of the top 1 percent; second, an estimate which includes an adjustment that tries to correct for this problem; and third, an estimate that includes an additional correction for hidden global wealth (assets held in tax havens).²² For the third estimate, we assume a rather strong (6 percent) return on the hidden assets, and we assume that all hidden assets belong to the global top 1 percent.²³ The income share of the richest 1 percent of people in 2010 increases from 15.7 percent under the first scenario, to 28 percent when we make an adjustment for top income underestimation in surveys, to 29 percent when we make an additional adjustment for income from hidden wealth. But all of these estimates of income share fall far short of the estimate of the global top 1 percent share in wealth made by the Credit Suisse Research Institute in 2013, which was 46 percent. From around 2000 to around 2010, the global income share of the top 1 percent either remained constant or increased slightly, while their global wealth share rose (Table 1.1).

There is thus a divergence in the evolution of income and wealth concentrations. According to the Credit Suisse Research Institute (2014), the increasing concentration in wealth is due to the strong performance of world stock markets after 2010 and to presumed higher rates of return received by the rich. The divergence between income and wealth concentrations for the top 1 percent is consistent with the picture of significant *income* gains realized by the middle of the global income distribution during the past thirty years. The growing incomes of this group have put something of a damper on the growth of the income share of the top 1 percent. But it is also very likely that the people around the global middle, who are still poor,

GLOBAL INEQUALITY

40

TABLE 1.1. Global top 1% shares in global income and global wealth

Estimate of income or wealth share	Around 2000	Around 2010
Top 1% share in global income based on household surveys alone ^a	14.5	15.7
Top 1% share in global income based on surveys and adjustment for	29	28
underreporting ^a Top 1% share in global income based on surveys, adjustment for underreporting,		29
and adjustment for hidden wealth ^b lop 1% share in global wealth ^c	32	46

Note: Top 1% for wealth refers to the richest 1% of adult individuals.

^a From Lakner and Milanovic (2013); methodology of imputation explained in the paper.

^b Additional data from Zucman (2013).

^c For 2000 from Davies et al. (2011, 244); for 2013 from Credit Suisse Research Institute (2013, 10, table 1).

have hardly any assets at all. Consequently, their asset growth must have been very small and could not have provided any offsetting effect to the rising amounts of wealth and thus wealth share of the top 1 percent.

The Real Global Plutocrats: The Billionaires

In 2013, according to the *Forbes* list of billionaires, there were 1,426 individuals in the world whose net worth was equal to or greater than \$1 billion.²⁴ This small and select group, together with their family members, represents one-hundredth of one-hundredth of the global 1 percent. Their total assets are estimated at \$5.4 trillion. According to a 2013 Credit Suisse report (p. 5, table 1), the world's wealth is estimated

EXCURSUS 1.3. What Is a Billion?

It is very difficult to comprehend what a number such as one billion really means. A billion dollars is so far outside the usual experience of practically everybody on earth that the very quantity it implies is not easily understood---other than that it is a very large amount indeed. It might help to think of it in the following manner. Suppose that a good fairy gave you one dollar each second. How much time would elapse before you collected \$1 million, and then \$1 billion? For the former, you would need 11.4 days; for the latter, almost thirty-two years. Or look at it from the consumption side. Suppose now that you inherited either \$1 million or \$1 billion, and that you spent \$1,000 every day. It would take you less than three years to run through your inheritance in the first case, and more than 2,700 years (that is, the time that separates us from Homer's *lliad*) to blow your inheritance in the second case. Or take the problem faced by drug lords. To transport \$1 million in \$100 bills requires a medium-sized briefcase. To ferry \$1 billion in the same banknotes would require a thousand such briefcases. Even if you used a big roller-bag, you would need about five hundred of them. And buying five hundred suitcases would attract attention that you might prefer to avoid.

at \$241 trillion. This means that this super-tiny group of individuals and their families controls about 2 percent of world wealth. To put it differently, these billionaires own twice as much wealth as exists in all of Africa.

How much has the wealth of the super-rich changed during globalization? *Forbes*'s annual lists give us a good approximate means to answer that question. It is important to realize, however, that in such lists, the cut-off point is an absolute level of wealth that gradu-

ally declines in real terms if there is inflation. Thus, a registered increase in the number of such individuals is in part spurious, due simply to the lowering of the real threshold. Methodologically, this "wealth line" is identical to the poverty line: in principle, we would like to fix the poverty (or wealth) line in real terms and then check to see if the number of individuals, or their share in the total population, has gone up or down. This is indeed what we routinely do for poverty lines. Here, we have to do the same for the wealth line. In order to fix the wealth line in real terms, we use the US Consumer Price Index (CPI). Very conveniently, it turns out that the wealth line of \$1 billion in 1987, when Forbes started publishing its global wealth lists, is equivalent in real terms to a wealth line of \$2 billion in 2013 (the US price index having exactly doubled over this period). For simplicity's sake, let's call the people above that constant real level (\$1 billion in 1987 prices) the hyper-wealthy or the hyper-rich.

Until 1992, *Forbes* published two separate lists: one of the four hundred richest Americans (which began in 1982), and another of global billionaires (started in 1987). In 1987, there were 49 billionaires in the United States and 96 billionaires in the rest of the world (thus in total there were 145 such individuals). *Forbes* did not calculate their combined wealth, but it may be estimated at \$450 billion.²⁵ These two numbers (145 hyper-wealthy people and \$450 billion) from 1987 are what we will use to compare with the number and wealth of bi-billionaires (that is, people with net wealth in excess of \$2 billion) in 2013. Conveniently, these two dates (1987 and 2013) bracket almost the same period from which we have household survey data (1988 to 2011) and thus allow us to look at what happened both on the income and the wealth sides.

In 2013, the number of bi-billionaires was 735, and their total wealth was \$4.5 trillion (equivalent to \$2.25 trillion in 1987 prices). Thus, both the number of hyper-wealthy people and their combined

real wealth have expanded by a factor of five (\$2.25 trillion versus \$0.45 trillion). An obvious implication of this rough calculation is that per capita wealth of billionaires has not gone up in real terms. The average wealth of the hyper-rich was about \$3 billion (in 1987 US prices) in both 1987 and 2013. There are simply many more of the hyper-rich now than there were in the late 1980s.

Meanwhile, the real world GDP has increased by 2.25 times, which is significantly less than the increase in the real wealth of the hyperrich. As a result, the share of the hyper-wealthy individuals expressed in terms of world GDP has more than doubled, from less than 3 percent to more than 6 percent (Figure 1.8).²⁶

These figures give us a reasonably firm grasp on the growth of the global plutocracy: their ranks, although tiny, have increased five-





This graph shows the total wealth of hyper-wealthy individuals as a share of global GDP. The hyper-wealthy are defined as people with net assets above \$1 billion in 1987 US prices (equal to \$2 billion in 2013 US prices). We see that their wealth increased from 1987 to 2013, relative to global GDP. Data source: Author's calculations from various *Forbes* lists. fold, and their total wealth, measured in terms of global GDP, has more doubled. This growth, together with the expansion of the emerging global middle class, is the most significant development of the high globalization era that began in the late 1980s. What these two developments-one that may be considered hopeful, and the other perhaps ominous-might imply for the coming decades will be explored in Chapter 4. First, however, we need to address an issue that we have so far barely mentioned: income inequalities within nations and their long-term evolution. That is the subject of Chapter 2. For global inequality, inequalities within nations do play a role, but today it is a subsidiary role because their influence on global inequality is less than the influence of differential growth rates of poor, middle-income, and rich countries. However, as we shall see in Chapter 3, this rather minor role of within-nation inequality has not always been the case and in the future might change again. Moreover, so far we have intentionally focused solely on changes in global magnitudes. But national inequalities are still the most important form of inequality from the political point of view. Our world is politically organized into nation-states, and it is inequalities within nations that people most frequently debate, on which they most ardently disagree, and on whose long-term movements there exist various theories. In the next chapter I discuss within-nation inequalities and propose an alternative theory of their long-run evolution that is more complete and satisfactory, in my view, than the existing theories.