CAPITAL IN THE TWENTY-FIRST CENTURY

Thomas Piketty

Translated by Arthur Goldhammer

The Belknap Press of Harvard University Press

CAMBRIDGE, MASSACHUSETTS LONDON, ENGLAND
2014

First published as *Le capital au XXI siècle,* copyright © 2013 Éditions du Seuil

Design by Dean Bornstein

Jacket design by Graciela Galup

The Library of Congress has cataloged the printed edition as follows
Piketty, Thomas, 1971–
[Capital au XXIe siècle. English]
Capital in the twenty-first century / Thomas Piketty ; translated by Arthur Goldhammer.
pages cm
Translation of the author's Le capital au XXIe siècle.
Includes bibliographical references and index.
ISBN 978-0-674-43000-6 (alk. paper)
1. Capital. 2. Income distribution. 3. Wealth. 4. Labor economics. I. Goldhammer, Arthur, translator. II. Title.
HB501.P43613 2014
332'.041—dc23
2013036024

{NINE}

Inequality of Labor Income

Now that I have introduced the evolution of income and wages in France and the United States since the beginning of the twentieth century, I will examine the changes I have observed and consider how representative they are of long-term changes in other developed and emerging economies.

I will begin by examining in this chapter the dynamics of labor income inequality. What caused the explosion of wage inequalities and the rise of the supermanager in the United States after 1980? More generally, what accounts for the diverse historical evolutions we see in various countries?

In subsequent chapters I will look into the evolution of the capital ownership distribution: How and why has the concentration of wealth decreased everywhere, but especially in Europe, since the turn of the twentieth century? The emergence of a "patrimonial middle class" is a crucial issue for this study, because it largely explains why income inequality decreased during the first half of the twentieth century and why we in the developed countries have gone from a society of rentiers to a society of managers (or, in the less optimistic version, from a society of superrentiers to a somewhat less extreme form of rentier society).

Wage Inequality: A Race between Education and Technology?

Why is inequality of income from labor, and especially wage inequality, greater in some societies and periods than others? The most widely accepted theory is that of a race between education and technology. To be blunt, this theory does not explain everything. In particular, it does not offer a satisfactory explanation of the rise of the supermanager or of wage inequality in the United States after 1980. The theory does, however, suggest interesting and important clues for explaining certain historical evolutions. I will therefore begin by discussing it.

The theory rests on two hypotheses. First, a worker's wage is equal to his marginal productivity, that is, his individual contribution to the output of the firm or office for which he works. Second, the worker's productivity depends above all on his skill and on supply and demand for that skill in a given society. For example, in a society in which very few people are qualified engineers (so that the "supply" of engineers is low) and the prevailing technology requires many engineers (so that "demand" is high), then it is highly likely that this combination of low supply and high demand will result in very high pay for engineers (relative to other workers) and therefore significant wage inequality between highly paid engineers and other workers.

This theory is in some respects limited and naïve. (In practice, a worker's productivity is not an immutable, objective quantity inscribed on his forehead, and the relative power of different social groups often plays a central role in determining what each worker is paid.) Nevertheless, as simple or even simplistic as the theory may be, it has the virtue of emphasizing two social and economic forces that do indeed play a fundamental role in determining wage inequality, even in more sophisticated theories: the supply and demand of skills. In practice, the supply of skills depends on, among other

things, the state of the educational system: how many people have access to this or that track, how good is the training, how much classroom teaching is supplemented by appropriate professional experience, and so on. The demand for skills depends on, among other things, the state of the technologies available to produce the goods and services that society consumes. No matter what other forces may be involved, it seems clear that these two factors—the state of the training system on the one hand, the state of technology on the other—play a crucial role. At a minimum, they influence the relative power of different social groups.

These two factors themselves depend on many other forces. The educational system is shaped by public policy, criteria of selection for different tracks, the way it is financed, the cost of study for students and their families, and the availability of continuing education. Technological progress depends on the pace of innovation and the rapidity of implementation. It generally increases the demand for new skills and creates new occupations. This leads to the idea of a race between education and technology: if the supply of skills does not increase at the same pace as the needs of technology, then groups whose training is not sufficiently advanced will earn less and be relegated to devalued lines of work, and inequality with respect to labor will increase. In order to avoid this, the educational system must increase its supply of new types of training and its output of new skills at a sufficiently rapid pace. If equality is to decrease, moreover, the supply of new skills must increase even more rapidly, especially for the least well educated.

Consider, for example, wage inequalities in France. As I have shown, the wage hierarchy was fairly stable over a long period of time. The average wage increased enormously over the course of the twentieth century, but the gap between the best and worst paid deciles remained the same. Why was this the case, despite the massive democratization of the educational system during the same period? The most natural explanation is that all skill levels progressed at roughly the same pace, so that the inequalities in the wage scale were simply translated upward. The bottom group, which had once only finished grade school, moved up a notch on the educational ladder, first completing junior high school, then going on to a high school diploma. But the group that had previously made do with a high school diploma now went on to college or even graduate school. In other words, the democratization of the educational system did not eliminate educational inequality and therefore did not reduce wage inequality. If educational democratization had not taken place, however, and if the children of those who had only finished grade school a century ago (three-quarters of each generation at that time) had remained at that level, inequalities with respect to labor, and especially wage inequalities, would surely have increased substantially.

Now consider the US case. Two economists, Claudia Goldin and Lawrence Katz, systematically compared the following two evolutions in the period 1890–2005: on the one hand the wage gap between workers who graduated from college and those who had only a high school diploma, and on the other the rate of growth of the number of college degrees. For Goldin and Katz, the conclusion is stark: the two curves move in opposite directions. In particular, the wage gap, which decreased fairly regularly until the 1970s, suddenly begins to widen in the 1980s, at precisely the moment when for the first time the number of college graduates stops growing, or at any rate grows much more slowly than before. Goldin and Katz have no doubt that increased wage inequality in the United States is due to a failure to invest sufficiently in higher education. More precisely, too many people failed to receive

the necessary training, in part because families could not afford the high cost of tuition. In order to reverse this trend, they conclude, the United States should invest heavily in education so that as many people as possible can attend college.

The lessons of French and US experience thus point in the same direction. In the long run, the best way to reduce inequalities with respect to labor as well as to increase the average productivity of the labor force and the overall growth of the economy is surely to invest in education. If the purchasing power of wages increased fivefold in a century, it was because the improved skills of the workforce, coupled with technological progress, increased output per head fivefold. Over the long run, education and technology are the decisive determinants of wage levels.

By the same token, if the United States (or France) invested more heavily in high-quality professional training and advanced educational opportunities and allowed broader segments of the population to have access to them, this would surely be the most effective way of increasing wages at the low to medium end of the scale and decreasing the upper decile's share of both wages and total income. All signs are that the Scandinavian countries, where wage inequality is more moderate than elsewhere, owe this result in large part to the fact that their educational system is relatively egalitarian and inclusive. The question of how to pay for education, and in particular how to pay for higher education, is everywhere one of the key issues of the twenty-first century. Unfortunately, the data available for addressing issues of educational cost and access in the United States and France are extremely limited. Both countries attach a great deal of importance to the central role of schools and vocational training in fostering social mobility, yet theoretical discussion of educational issues and of meritocracy is often out of touch with reality, and in particular with the fact that the most prestigious schools tend to favor students from privileged social backgrounds. I will come back to this point in Chapter 13.

The Limits of the Theoretical Model: The Role of Institutions

Education and technology definitely play a crucial role in the long run. This theoretical model, based on the idea that a worker's wage is always perfectly determined by her marginal productivity and thus primarily by skill, is nevertheless limited in a number of ways. Leave aside the fact that it is not always enough to invest in training: existing technology is sometimes unable to make use of the available supply of skills. Leave aside, too, the fact that this theoretical model, at least in its most simplistic form, embodies a far too instrumental and utilitarian view of training. The main purpose of the health sector is not to provide other sectors with workers in good health. By the same token, the main purpose of the educational sector is not to prepare students to take up an occupation in some other sector of the economy. In all human societies, health and education have an intrinsic value: the ability to enjoy years of good health, like the ability to acquire knowledge and culture, is one of the fundamental purposes of civilization. We are free to imagine an ideal society in which all other tasks are almost totally automated and each individual has as much freedom as possible to pursue the goods of education, culture, and health for the benefit of herself and others. Everyone would be by turns teacher or student, writer or reader, actor or spectator, doctor or patient. As noted in Chapter 2, we are to some extent already on this path: a characteristic feature of modern growth is the considerable share of both output and employment devoted to education, culture, and medicine.

While awaiting the ideal society of the future, let us try to gain a better understanding of wage inequality today. In this narrower context, the main problem with the theory of marginal productivity is quite simply that it fails to explain the diversity of the wage distributions we observe in different countries at different times. In order to understand the dynamics of wage inequality, we must introduce other factors, such as the institutions and rules that govern the operation of the labor market in each society. To an even greater extent than other markets, the labor market is not a mathematical abstraction whose workings are entirely determined by natural and immutable mechanisms and implacable technological forces: it is a social construct based on specific rules and compromises.

In the previous chapter I noted several important episodes of compression and expansion of wage hierarchies that are very difficult to explain solely in terms of the supply of and demand for various skills. For example, the compression of wage inequalities that occurred in both France and the United States during World Wars I and II was the result of negotiations over wage scales in both the public and private sectors, in which specific institutions such as the National War Labor Board (created expressly for the purpose) played a central role. I also called attention to the importance of changes in the minimum wage for explaining the evolution of wage inequalities in France since 1950, with three clearly identified subperiods: 1950–1968, during which the minimum wage was rarely adjusted and the wage hierarchy expanded; 1968–1983, during which the minimum wage rose very rapidly and wage inequalities decreased sharply; and finally 1983–2012, during which the minimum wage increased relatively slowly and the wage hierarchy tended to expand. At the beginning of 2013, the minimum wage in France stood at 9.43 euros per hour.

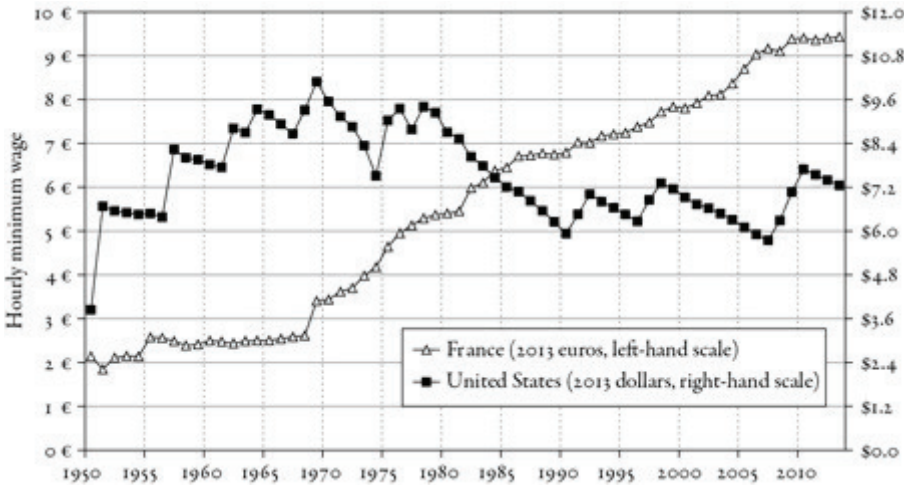

FIGURE 9.1. Minimum wage in France and the United States, 1950–2013

Expressed in 2013 purchasing power, the hourly minimum wage rose from $3.80 to $7.30 between 1950 and 2013 in the United States, and from €2.10 to €9.40 in France.

Sources and series: see piketty.pse.ens.fr/capital21c.

In the United States, a federal minimum wage was introduced in 1933, nearly twenty years earlier than in France. As in France, changes in the minimum wage played an important role in the evolution of wage inequalities in the United States. It is striking to learn that in terms of purchasing power, the minimum wage reached its maximum level nearly half a century ago, in 1969, at $1.60 an hour (or $10.10 in 2013 dollars, taking account of inflation between 1968 and 2013), at a time when the unemployment rate was below 4 percent. From 1980 to 1990, under the presidents Ronald Reagan and George H. W. Bush, the federal minimum wage remained stuck at $3.35, which led to a

significant decrease in purchasing power when inflation is factored in. It then rose to $5.25 under Bill Clinton in the 1990s and was frozen at that level under George W. Bush before being increased several times by Barack Obama after 2008. At the beginning of 2013 it stood at $7.25 an hour, or barely 6 euros, which is a third below the French minimum wage, the opposite of the situation that obtained in the early 1980s (see Figure 9.1). President Obama, in his State of the Union address in February 2013, announced his intention to raise the minimum wage to about $9 an hour for the period 2013–2016.

Inequalities at the bottom of the US wage distribution have closely followed the evolution of the minimum wage: the gap between the bottom 10 percent of the wage distribution and the overall average wage widened significantly in the 1980s, then narrowed in the 1990s, and finally increased again in the 2000s. Nevertheless, inequalities at the top of the distribution—for example, the share of total wages going to the top 10 percent—increased steadily throughout this period. Clearly, the minimum wage has an impact at the bottom of the distribution but much less influence at the top, where other forces are at work.

Wage Scales and the Minimum Wage

There is no doubt that the minimum wage plays an essential role in the formation and evolution of wage inequalities, as the French and US experiences show. Each country has its own history in this regard and its own peculiar chronology. There is nothing surprising about that: labor market regulations depend on each society's perceptions and norms of social justice and are intimately related to each country's social, political, and cultural history. The United States used the minimum wage to increase lower-end wages in the 1950s and 1960s but abandoned this tool in the 1970s. In France, it was exactly the opposite: the minimum wage was frozen in the 1950s and 1960s but was used much more often in the 1970s. Figure 9.1 illustrates this striking contrast.

It would be easy to multiply examples from other countries. Britain introduced a minimum wage in 1999, at a level between the United States and France: in 2013 it was £6.19 (or about 8.05 euros). Germany and Sweden have chosen to do without minimum wages at the national level, leaving it to trade unions to negotiate not only minimums but also complete wage schedules with employers in each branch of industry. In practice, the minimum wage in both countries was about 10 euros an hour in 2013 in many branches (and therefore higher than in countries with a national minimum wage). But minimum pay can be markedly lower in sectors that are relatively unregulated or underunionized. In order to set a common floor, Germany is contemplating the introduction of a minimum wage in 2013–2014. This is not the place to write a detailed history of minimum wages and wage schedules around the world or to discuss their impact on wage inequality. My goal here is more modest: simply to indicate briefly what general principles can be used to analyze the institutions that regulate wage setting everywhere.

What is in fact the justification for minimum wages and rigid wage schedules? First, it is not always easy to measure the marginal productivity of a particular worker. In the public sector, this is obvious, but it is also clear in the private sector: in an organization employing dozens or even thousands of workers, it is no simple task to judge each individual worker's contribution to overall output. To be sure, one can estimate marginal productivity, at least for jobs that can be replicated, that

is, performed in the same way by any number of employees. For an assembly-line worker or McDonald's server, management can calculate how much additional revenue an additional worker or server would generate. Such an estimate would be approximate, however, yielding a range of productivities rather than an absolute number. In view of this uncertainty, how should the wage be set? There are many reasons to think that granting management absolute power to set the wage of each employee on a monthly or (why not?) daily basis would not only introduce an element of arbitrariness and injustice but would also be inefficient for the firm.

In particular, it may be efficient for the firm to ensure that wages remain relatively stable and do not vary constantly with fluctuations in sales. The owners and managers of the firm usually earn much more and are significantly wealthier than their workers and can therefore more easily absorb short-term shocks to their income. Under such circumstances, it can be in everyone's interest to provide a kind of "wage insurance" as part of the employment contract, in the sense that the worker's monthly wage is guaranteed (which does not preclude the use of bonuses and other incentives). The payment of a monthly rather than a daily wage was a revolutionary innovation that gradually took hold in all the developed countries during the twentieth century. This innovation was inscribed in law and became a feature of wage negotiations between workers and employers. The daily wage, which had been the norm in the nineteenth century, gradually disappeared. This was a crucial step in the constitution of the working class: workers now enjoyed a legal status and received a stable, predictable remuneration for their work. This clearly distinguished them from day laborers and piece workers—the typical employees of the eighteenth and nineteenth centuries.

This justification of setting wages in advance obviously has its limits. The other classic argument in favor of minimum wages and fixed wage schedules is the problem of "specific investments." Concretely, the particular functions and tasks that a firm needs to be performed often require workers to make specific investments in the firm, in the sense that these investments are of no (or limited) value to other firms: for instance, workers might need to learn specific work methods, organizational methods, or skills linked to the firm's production process. If wages can be set unilaterally and changed at any moment by the firm, so that workers do not know in advance how much they will be paid, then it is highly likely that they will not invest as much in the firm as they should. It may therefore be in everyone's interest to set pay scales in advance. The same "specific investments" argument can also apply to other decisions by the firm, and it is the main reason for limiting the power of stockholders (who are seen as having too short-term an outlook in some cases) in favor of a power-sharing arrangement with a broader group of "stakeholders" (including the firm's workers), as in the "Rhenish model" of capitalism discussed earlier, in Part Two. This is probably the most important argument in favor of fixed wage scales.

More generally, insofar as employers have more bargaining power than workers and the conditions of "pure and perfect" competition that one finds in the simplest economic models fail to be satisfied, it may be reasonable to limit the power of employers by imposing strict rules on wages. For example, if a small group of employers occupies a monopsony position in a local labor market (meaning that they are virtually the only source of employment, perhaps because of the limited mobility of the local labor force), they will probably try to exploit their advantage by lowering wages as much as possible, possibly even below the marginal productivity of the workers. Under such conditions, imposing a

minimum wage may be not only just but also efficient, in the sense that the increase in wages may move the economy closer to the competitive equilibrium and increase the level of employment. This theoretical model, based on imperfect competition, is the clearest justification for the existence of a minimum wage: the goal is to make sure that no employer can exploit his competitive advantage beyond a certain limit.

Again, everything obviously depends on the level of the minimum wage. The limit cannot be set in the abstract, independent of the country's general skill level and average productivity. Various studies carried out in the United States between 1980 and 2000, most notably by the economists David Card and Alan Krueger, showed that the US minimum wage had fallen to a level so low in that period that it could be raised without loss of employment, indeed at times with an increase in employment, as in the monopsony model. On the basis of these studies, it seems likely that the increase in the minimum wage of nearly 25 percent (from $7.25 to $9 an hour) currently envisaged by the Obama administration will have little or no effect on the number of jobs. Obviously, raising the minimum wage cannot continue indefinitely: as the minimum wage increases, the negative effects on the level of employment eventually win out. If the minimum wage were doubled or tripled, it would be surprising if the negative impact were not dominant. It is more difficult to justify a significant increase in the minimum wage in a country like France, where it is relatively high (compared with the average wage and marginal productivity), than in the United States. To increase the purchasing power of low-paid workers in France, it is better to use other tools, such as training to improve skills or tax reform (these two remedies are complementary, moreover). Nevertheless, the minimum wage should not be frozen. Wage increases cannot exceed productivity increases indefinitely, but it is just as unhealthy to restrain (most) wage increases to below the rate of productivity increase. Different labor market institutions and policies play different roles, and each must be used in an appropriate manner.

To sum up: the best way to increase wages and reduce wage inequalities in the long run is to invest in education and skills. Over the long run, minimum wages and wage schedules cannot multiply wages by factors of five or ten: to achieve that level of progress, education and technology are the decisive forces. Nevertheless, the rules of the labor market play a crucial role in wage setting during periods of time determined by the relative progress of education and technology. In practice, those periods can be fairly long, in part because it is hard to gauge individual marginal productivities with any certainty, and in part because of the problem of specific investments and imperfect competition.

How to Explain the Explosion of Inequality in the United States?

The most striking failure of the theory of marginal productivity and the race between education and technology is no doubt its inability to adequately explain the explosion of very high incomes from labor observed in the United States since 1980. According to this theory, one should be able to explain this change as the result of skill-biased technological change. Some US economists buy this argument, which holds that top labor incomes have risen much more rapidly than average wages simply because unique skills and new technology have made these workers much more productive than the average. There is a certain tautological quality to this explanation (after all, one can "explain" any distortion of the wage hierarchy as the result of some supposed technological change). It also has other major weaknesses, which to my mind make it a rather unconvincing argument.

First, as shown in the previous chapter, the increase in wage inequality in the United States is due mainly to increased pay at the very top end of the distribution: the top 1 percent and even more the top 0.1 percent. If we look at the entire top decile, we find that "the 9 percent" have progressed more rapidly than the average worker but not nearly at the same rate as "the 1 percent." Concretely, those making between $100,000 and $200,000 a year have seen their pay increase only slightly more rapidly than the average, whereas those making more than $500,000 a year have seen their remuneration literally explode (and those above $1 million a year have risen even more rapidly). This very sharp discontinuity at the top income levels is a problem for the theory of marginal productivity: when we look at the changes in the skill levels of different groups in the income distribution, it is hard to see any discontinuity between "the 9 percent" and "the 1 percent," regardless of what criteria we use: years of education, selectivity of educational institution, or professional experience. One would expect a theory based on "objective" measures of skill and productivity to show relatively uniform pay increases within the top decile, or at any rate increases within different subgroups much closer to one another than the widely divergent increases we observe in practice.

Make no mistake: I am not denying the decisive importance of the investments in higher education and training that Katz and Goldin have identified. Policies to encourage broader access to universities are indispensable and crucial in the long run, in the United States and elsewhere. As desirable as such policies are, however, they seem to have had limited impact on the explosion of the topmost incomes observed in the United States since 1980.

In short, two distinct phenomena have been at work in recent decades. First, the wage gap between college graduates and those who go no further than high school has increased, as Goldin and Katz showed. In addition, the top 1 percent (and even more the top 0.1 percent) have seen their remuneration take off. This is a very specific phenomenon, which occurs within the group of college graduates and in many cases separates individuals who have pursued their studies at elite universities for many years. Quantitatively, the second phenomenon is more important than the first. In particular, as shown in the previous chapter, the overperformance of the top centile explains most (nearly three-quarters) of the increase in the top decile's share of US national income since 1970. It is therefore important to find an adequate explanation of this phenomenon, and at first sight the educational factor does not seem to be the right one to focus on.

The Rise of the Supermanager: An Anglo-Saxon Phenomenon

The second difficulty—and no doubt the major problem confronting the marginal productivity theory—is that the explosion of very high salaries occurred in some developed countries but not others. This suggests that institutional differences between countries rather than general and a priori universal causes such as technological change played a central role.

I begin with the English-speaking countries. Broadly speaking, the rise of the supermanager is largely an Anglo-Saxon phenomenon. Since 1980 the share of the upper centile in national income has risen significantly in the United States, Great Britain, Canada, and Australia (see Figure 9.2). Unfortunately, we do not have separate series for wage inequality and total income inequality for all countries as we do for France and the United States. But in most cases we do have data concerning the composition of income in relation to total income, from which we can infer that in all of these

countries the explosion of top incomes explains most (generally at least two-thirds) of the increase in the top centile's share of national income; the rest is explained by robust income from capital. In all the English-speaking countries, the primary reason for increased income inequality in recent decades is the rise of the supermanager in both the financial and nonfinancial sectors.

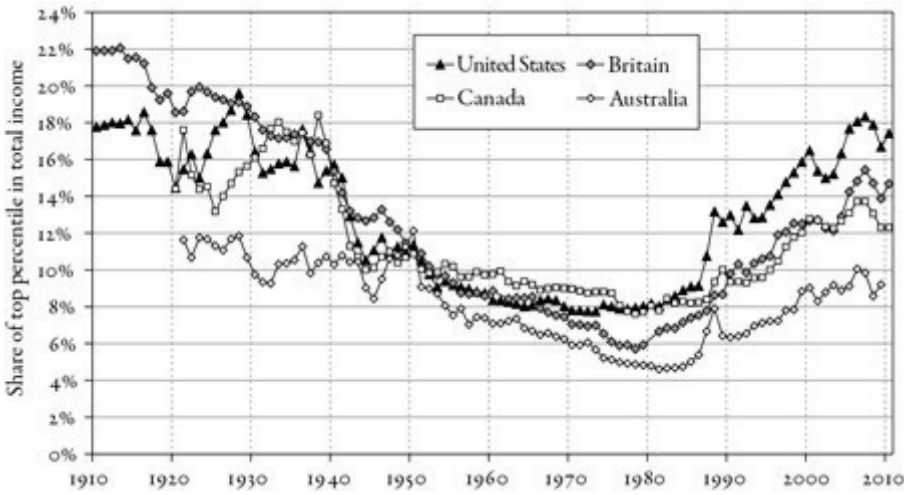

FIGURE 9.2. Income inequality in Anglo-Saxon countries, 1910–2010
The share of top percentile in total income rose since the 1970s in all Anglo-Saxon countries, but with different magnitudes.
Sources and series: see piketty.pse.ens.fr/capital21c.

This family resemblance should not be allowed to obscure the fact that the magnitude of the phenomenon varies widely from country to country, however. Figure 9.2 is quite clear on this point. In the 1970s, the upper centile's share of national income was quite similar across countries. It ranged from 6 to 8 percent in the four English-speaking countries considered, and the United States did not stand out as exceptional: indeed, Canada was slightly higher, at 9 percent, whereas Australia came in last, with just 5 percent of national income going to the top centile in the late 1970s and early 1980s. Thirty years later, in the early 2010s, the situation is totally different. The upper centile's share is nearly 20 percent in the United States, compared with 14–15 percent in Britain and Canada and barely 9–10 percent in Australia (see Figure 9.2). To a first approximation, we can say that the upper centile's share in the United States increased roughly twice as much as in Britain and Canada and about three times as much as in Australia and New Zealand. If the rise of the supermanager were a purely technological phenomenon, it would be difficult to understand why such large differences exist between otherwise quite similar countries.

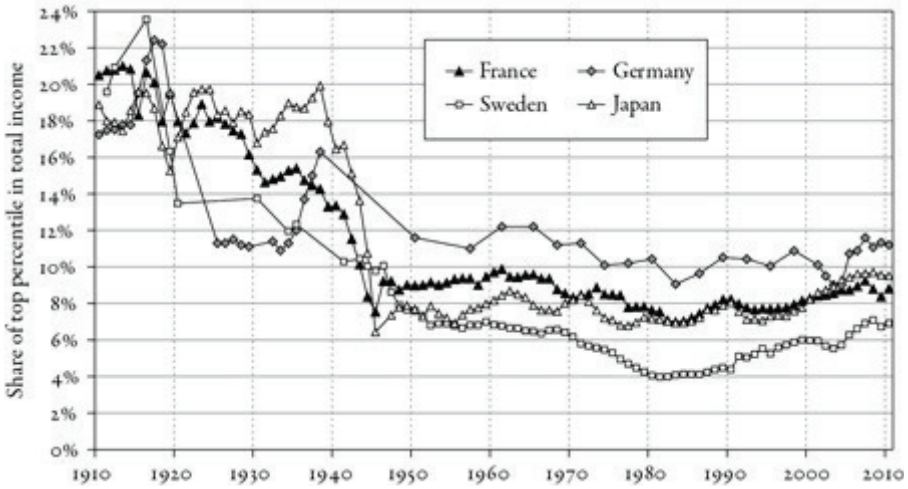

FIGURE 9.3. Income inequality in Continental Europe and Japan, 1910–2010

As compared to Anglo-Saxon countries, the share of top percentile barely increased since the 1970s in Continental Europe and Japan.

Sources and series: see piketty.pse.ens.fr/capital21c.

Let me turn now to the rest of the wealthy world, namely, continental Europe and Japan. The key fact is that the upper centile's share of national income in these countries has increased much less than in the English-speaking countries since 1980. The comparison between Figures 9.2 and 9.3 is particularly striking. To be sure, the upper centile's share increased significantly everywhere. In Japan the evolution was virtually the same as in France: the top centile's share of national income was barely 7 percent in the 1980s but is 9 percent or perhaps even slightly higher today. In Sweden, the top centile's share was a little more than 4 percent in the early 1980s (the lowest level recorded in the World Top Incomes Database for any country in any period) but reached 7 percent in the early 2010s. In Germany, the top centile's share rose from about 9 percent to nearly 11 percent of national income between the early 1980s and the early 2010s (see Figure 9.3).

If we look at other European countries, we observe similar evolutions, with the top centile's share increasing by two or three points of national income over the past thirty years in both northern and southern Europe. In Denmark and other Nordic countries, top incomes claim a smaller share of the total, but the increase is similar: the top centile received a little more than 5 percent of Danish national income in the 1980s but got close to 7 percent in 2000–2010. In Italy and Spain, the orders of magnitude are very close to those observed in France, with the top centile's share rising from 7 to 9 percent of national income in the same period, again an increase of two points of national income (see Figure 9.4). In this respect, continental Europe is indeed an almost perfect "union." Britain, of course, stands apart, being much closer to the pattern of the United States than that of Europe.

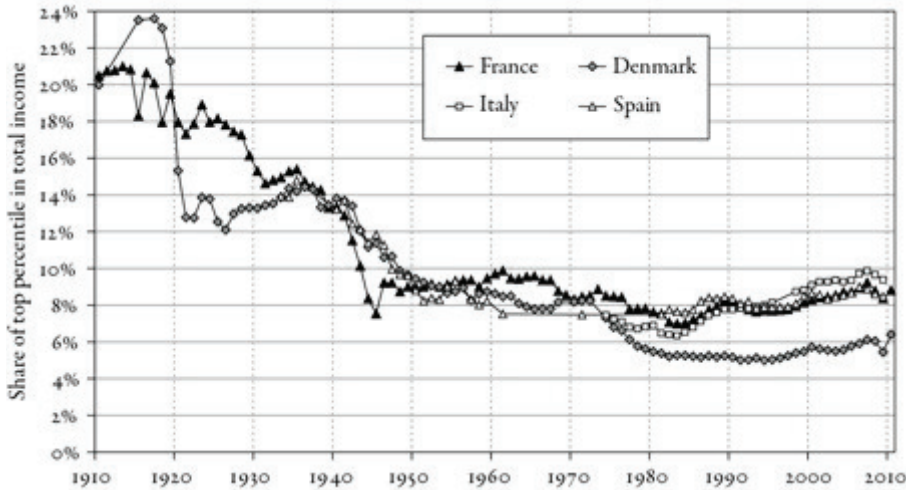

FIGURE 9.4. Income inequality in Northern and Southern Europe, 1910–2010

As compared to Anglo-Saxon countries, the top percentile income share barely increased in Northern and Southern Europe since the 1970s.

Sources and series: see piketty.pse.ens.fr/capital21c.

Make no mistake: these increases on the order of two to three points of national income in Japan and the countries of continental Europe mean that income inequality rose quite significantly. The top 1 percent of earners saw pay increases noticeably more rapid than the average: the upper centile's share increased by about 30 percent, and even more in countries where it started out lower. This was quite striking to contemporary observers, who read in the daily paper or heard on the radio about stupendous raises for "supermanagers." It was particularly striking in the period 1990–2010, when average income stagnated, or at least rose much more slowly than in the past.

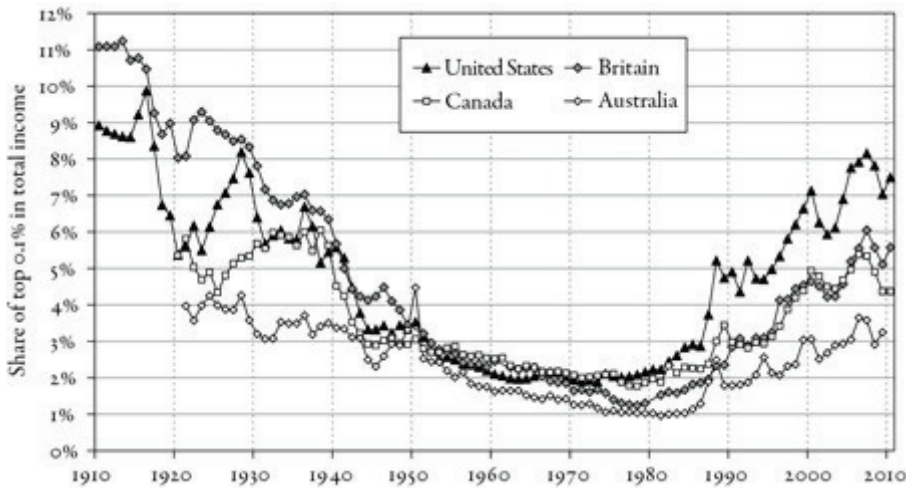

FIGURE 9.5. The top decile income share in Anglo-Saxon countries, 1910–2010

The share of the top 0.1 percent highest incomes in total income rose sharply since the 1970s in all Anglo-Saxon countries, but with varying magnitudes.

Sources and series: see piketty.pse.ens.fr/capital21c.

Furthermore, the higher one climbs in the income hierarchy, the more spectacular the raises. Even if the number of individuals benefiting from such salary increases is fairly limited, they are nevertheless quite visible, and this visibility naturally raises the question of what justifies such high levels of compensation. Consider the share of the top thousandth—the best remunerated 0.1 percent—in the national income of the English-speaking countries on the one hand (Figure 9.5) and continental Europe

and Japan on the other (Figure 9.6). The differences are obvious: the top thousandth in the United States increased their share from 2 to nearly 10 percent over the past several decades—an unprecedented rise. But there has been a remarkable increase of top incomes everywhere. In France and Japan, the top thousandth's share rose from barely 1.5 percent of national income in the early 1980s to nearly 2.5 percent in the early 2010s—close to double. In Sweden, the same share rose from less than 1 percent to more than 2 percent in the same period.

To make clear what this represents in concrete terms, remember that a 2 percent share of national income for 0.1 percent of the population means that the average individual in this group enjoys an income 20 times higher than the national average (or 600,000 euros a year if the average income is 30,000 per adult). A share of 10 percent means that each individual enjoys an income 100 times the national average (or 3 million euros a year if the average is 30,000). Recall, too, that the top 0.1 percent is by definition a group of 50,000 people in a country with a population of 50 million adults (like France in the early 2010s). This is a very small minority ("the 1 percent" is of course 10 times larger), yet it occupies a significant place in the social and political landscape. The central fact is that in all the wealthy countries, including continental Europe and Japan, the top thousandth enjoyed spectacular increases in purchasing power in 1990–2010, while the average person's purchasing power stagnated.

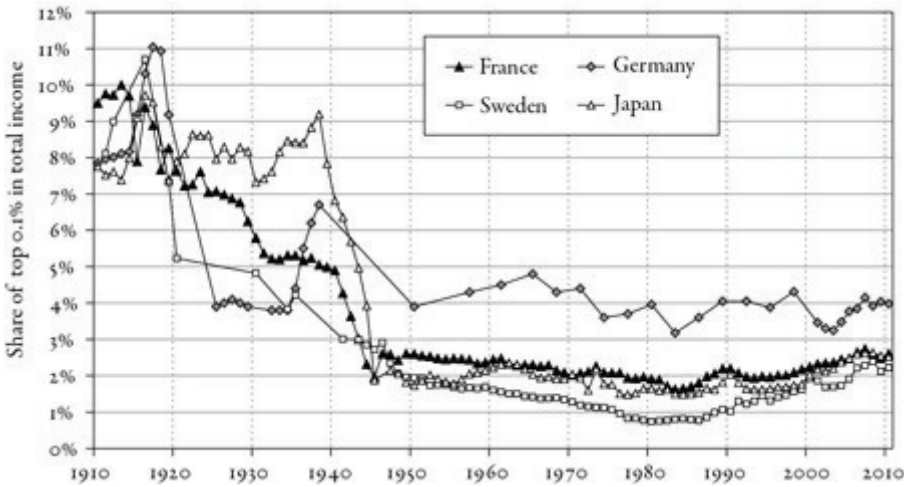

FIGURE 9.6. The top decile income share in Continental Europe and Japan, 1910–2010
As compared to Anglo-Saxon countries, the top 0.1 percent income share barely increased in Continental Europe and Japan.
Sources and series: see piketty.pse.ens.fr/capital21c.

From a macroeconomic point of view, however, the explosion of very high incomes has thus far been of limited importance in continental Europe and Japan: the rise has been impressive, to be sure, but too few people have been affected to have had an impact as powerful as in the United States. The transfer of income to "the 1 percent" involves only two to three points of national income in continental Europe and Japan compared with 10 to 15 points in the United States—5 to 7 times greater.

The simplest way to express these regional differences is no doubt the following: in the United States, income inequality in 2000–2010 regained the record levels observed in 1910–1920 (although the composition of income was now different, with a larger role played by high incomes from labor and a smaller role by high incomes from capital). In Britain and Canada, things moved in the same

direction. In continental Europe and Japan, income inequality today remains far lower than it was at the beginning of the twentieth century and in fact has not changed much since 1945, if we take a long-run view. The comparison of Figures 9.2 and 9.3 is particularly clear on this point.

Obviously, this does not mean that the European and Japanese evolutions of the past few decades should be neglected. On the contrary: their trajectory resembles that of the United States in some respects, with a delay of one or two decades, and one need not wait until the phenomenon assumes the macroeconomic significance observed in the United States to worry about it.

Nevertheless, the fact remains that the evolution in continental Europe and Japan is thus far much less serious than in the United States (and, to a lesser extent, in the other Anglo-Saxon countries). This may tell us something about the forces at work. The divergence between the various regions of the wealthy world is all the more striking because technological change has been the same more or less everywhere: in particular, the revolution in information technology has affected Japan, Germany, France, Sweden, and Denmark as much as the United States, Britain, and Canada. Similarly, economic growth—or, more precisely, growth in output per capita, which is to say, productivity growth—has been quite similar throughout the wealthy countries, with differences of a few tenths of a percentage point. In view of these facts, this quite large divergence in the way the income distribution has evolved in the various wealthy countries demands an explanation, which the theory of marginal productivity and of the race between technology and education does not seem capable of providing.

Europe: More Inegalitarian Than the New World in 1900–1910

Note, moreover, that the United States, contrary to what many people think today, was not always more inegalitarian than Europe—far from it. Income inequality was actually quite high in Europe at the beginning of the twentieth century. This is confirmed by all the indices and historical sources. In particular, the top centile's share of national income exceeded 20 percent in all the countries of Europe in 1900–1910 (see Figures 9.2–4). This was true not only of Britain, France, and Germany but also of Sweden and Denmark (proof that the Nordic countries have not always been models of equality—far from it), and more generally of all European countries for which we have estimates from this period.

The similar levels of income concentration in all European countries during the Belle Époque obviously demand an explanation. Since top incomes in this period consisted almost entirely of income from capital, the explanation must be sought primarily in the realm of concentration of capital. Why was capital so concentrated in Europe in the period 1900–1910?

It is interesting to note that, compared with Europe, inequality was lower not only in the United States and Canada (where the top centile's share of national income was roughly 16–18 percent at the beginning of the twentieth century) but especially in Australia and New Zealand (11–12 percent). Thus it was the New World, and especially the newest and most recently settled parts of the New World, that appear to have been less inegalitarian than Old Europe in the Belle Époque.

It is also interesting to note that Japan, despite its social and cultural differences from Europe, seems to have had the same high level of inequality at the beginning of the twentieth century, without about 20 percent of national income going to the top centile. The available data do not allow me to

make all the comparisons I would like to make, but all signs are that in terms of both income structure and income inequality, Japan was indeed part of the same "old world" as Europe. It is also striking to note the similar evolution of Japan and Europe over the course of the twentieth century (Figure 9.3).

I will return later to the reasons for the very high concentration of capital in the Belle Époque and to the transformations that took place in various countries over the course of the twentieth century (namely, a reduction of concentration). I will show in particular that the greater inequality of wealth that we see in Europe and Japan is fairly naturally explained by the low demographic growth rate we find in the Old World, which resulted almost automatically in a greater accumulation and concentration of capital.

At this stage, I want simply to stress the magnitude of the changes that have altered the relative standing of countries and continents. The clearest way to make this point is probably to look at the evolution of the top decile's share of national income. Figure 9.7 shows this for the United States and four European countries (Britain, France, Germany, and Sweden) since the turn of the twentieth century. I have indicated decennial averages in order to focus attention on long-term trends.

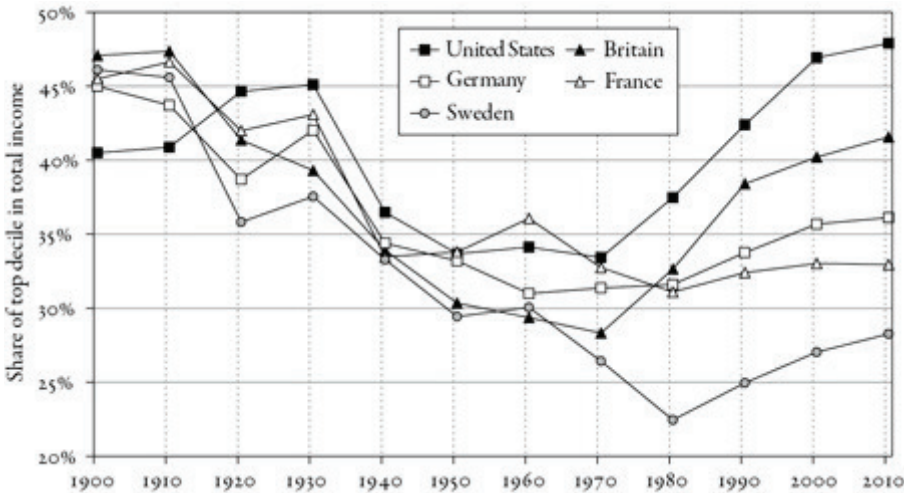

FIGURE 9.7. The top decile income share in Europe and the United States, 1900–2010
In the 1950s–1970s, the top decile income share was about 30–35 percent of total income in Europe as in the United States.
Sources and series: see piketty.pse.ens.fr/capital21c.

What we find is that on the eve of World War I, the top decile's share was 45–50 percent of national income in all the European countries, compared with a little more than 40 percent in the United States. By the end of World War II, the United States had become slightly more inegalitarian than Europe: the top decile's share decreased on both continents owing to the shocks of 1914–1945, but the fall was more precipitous in Europe (and Japan). The explanation for this is that the shocks to capital were much larger. Between 1950 and 1970, the upper decile's share was fairly stable and fairly similar in the United States and Europe, around 30–35 percent of national income. The strong divergence that began in 1970–1980 led to the following situation in 2000–2010: the top decile's share of US national income reached 45–50 percent, or roughly the same level as Europe in 1900–1910. In Europe, we see wide variation, from the most inegalitarian case (Britain, with a top decile share of 40 percent) to the most egalitarian (Sweden, less than 30 percent), with France and Germany in between (around 35 percent).

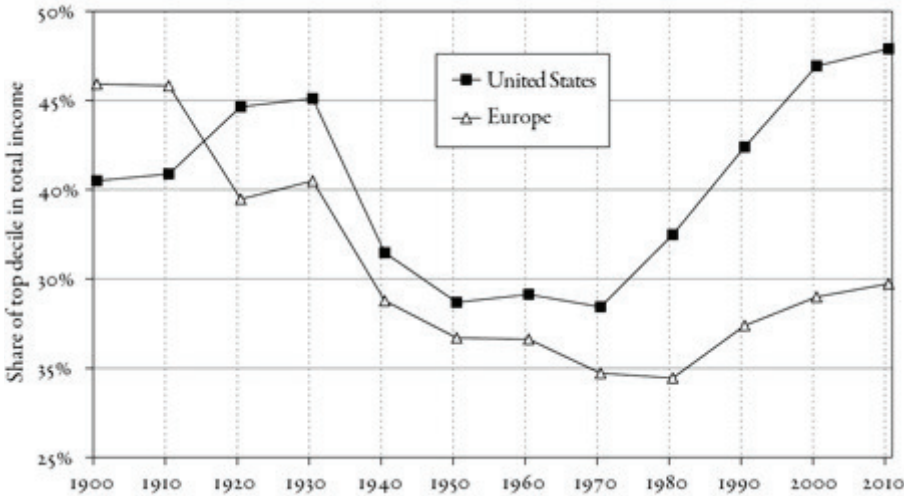

FIGURE 9.8. Income inequality in Europe versus the United States, 1900–2010

The top decile income share was higher in Europe than in the United States in 1900–1910; it is a lot higher in the United States in 2000–2010.

Sources and series: see piketty.pse.ens.fr/capital21c.

If we calculate (somewhat abusively) an average for Europe based on these four countries, we can make a very clear international comparison: the United States was less inegalitarian than Europe in 1900–1910, slightly more inegalitarian in 1950–1960, and much more inegalitarian in 2000–2010 (see Figure 9.8).

Apart from this long-term picture, there are of course multiple national histories as well as constant short- and medium-term fluctuations linked to social and political developments in each country, as I showed in Chapter 8 and analyzed in some detail in the French and US cases. Space will not permit me to do the same for every country here.

In passing, however, it is worth mentioning that the period between the two world wars seems to have been particularly tumultuous and chaotic almost everywhere, though the chronology of events varied from country to country. In Germany, the hyperinflation of the 1920s followed hard on the heels of military defeat. The Nazis came to power a short while later, after the worldwide depression had plunged the country back into crisis. Interestingly, the top centile's share of German national income increased rapidly between 1933 and 1938, totally out of phase with other countries: this reflects the revival of industrial profits (boosted by demand for armaments), as well as a general reestablishment of income hierarchies in the Nazi era. Note, too, that the share of the top centile—and, even more, the top thousandth—in Germany has been noticeably higher since 1950 than in most other continental European countries (including, in particular, France) as well as Japan, even though the overall level of inequality in Germany is not very different. This can be explained in various ways, among which it is difficult to say that one is better than another. (I will come back to this point.)

In addition, there are serious lacunae in German tax records, owing in large part to the country's turbulent history in the twentieth century, so that it is difficult to be sure about certain developments or to make sharp comparisons with other countries. Prussia, Saxony, and most other German states imposed an income tax relatively early, between 1880 and 1890, but there were no national laws or tax records until after World War I. There were frequent breaks in the statistical record during the 1920s, and then the records for 1938 to 1950 are missing altogether, so it is impossible to study how the income distribution evolved during World War II and its immediate aftermath.

This distinguishes Germany from other countries deeply involved in the conflict, especially Japan and France, whose tax administrations continued to compile statistics during the war years without interruption, as if nothing were amiss. If Germany was anything like these two countries, it is likely that the top centile's share of national income reached a nadir in 1945 (the year in which German capital and income from capital were reduced to virtually nothing) before beginning to rise sharply again in 1946–1947. Yet when German tax records return in 1950, they show the income hierarchy already beginning to resemble its appearance in 1938. In the absence of complete sources, it is difficult to say more. The German case is further complicated by the fact that the country's boundaries changed several times during the twentieth century, most recently with the reunification of 1990–1991, in addition to which full tax data are published only every three years (rather than annually as in most other countries).

Inequalities in Emerging Economies: Lower Than in the United States?

Let me turn now to the poor and emerging economies. The historical sources we need in order to study the long-run dynamics of the wealth distribution there are unfortunately harder to come by than in the rich countries. There are, however, a number of poor and emerging economies for which it is possible to find long series of tax data useful for making (rough) comparisons with our results for the more developed economies. Shortly after Britain introduced a progressive income tax at home, it decided to do the same in a number of its colonies. Thus an income tax fairly similar to that introduced in Britain in 1909 was adopted in South Africa in 1913 and in India (including present-day Pakistan) in 1922. Similarly, the Netherlands imposed an income tax on its Indonesian colony in 1920. Several South American countries introduced an income tax between the two world wars: Argentina, for example, did so in 1932. For these four countries—South Africa, India, Indonesia, and Argentina—we have tax data going back, respectively, to 1913, 1922, 1920, and 1932 and continuing (with gaps) to the present. The data are similar to what we have for the rich countries and can be employed using similar methods, in particular national income estimates for each country going back to the turn of the twentieth century.

My estimates are indicated in Figure 9.9. Several points deserve to be emphasized. First, the most striking result is probably that the upper centile's share of national income in poor and emerging economies is roughly the same as in the rich economies. During the most inegalitarian phases, especially 1910–1950, the top centile took around 20 percent of national income in all four countries: 15–18 percent in India and 22–25 percent in South Africa, Indonesia, and Argentina. During more egalitarian phases (essentially 1950–1980), the top centile's share fell to between 6 and 12 percent (barely 5–6 percent in India, 8–9 percent in Indonesia and Argentina, and 11–12 percent in South Africa). Thereafter, in the 1980s, the top centile's share rebounded, and today it stands at about 15 percent of national income (12–13 percent in India and Indonesia and 16–18 percent in South Africa and Argentina).

Figure 9.9 also shows two countries for which the available tax records allow us only to study how things have changed since the mid-1980s: China and Colombia. In China, the top centile's share of national income rose rapidly over the past several decades but starting from a fairly low (almost Scandinavian) level in the mid-1980s: less than 5 percent of national income went to the top centile at

that time, according to the available sources. This is not very surprising for a Communist country with a very compressed wage schedule and virtual absence of private capital. Chinese inequality increased very rapidly following the liberalization of the economy in the 1980s and accelerated growth in the period 1990–2000, but according to my estimates, the upper centile's share in 2000–2010 was 10–11 percent, less than in India or Indonesia (12–14 percent, roughly the same as Britain and Canada) and much lower than in South Africa or Argentina (16–18 percent, approximately the same as the United States).

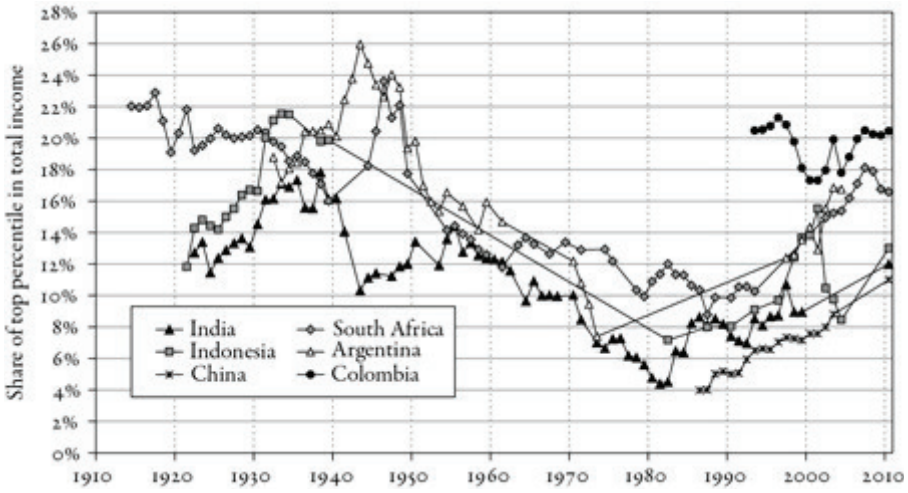

FIGURE 9.9. Income inequality in emerging countries, 1910–2010

Measured by the top percentile income share, income inequality rose in emerging countries since the 1980s, but ranks below the US level in 2000–2010.

Sources and series: see piketty.pse.ens.fr/capital21c.

Colombia on the other hand is one of the most inegalitarian societies in the WTID: the top centile's share stood at about 20 percent of national income throughout the period 1990–2010, with no clear trend (see Figure 9.9). This level of inequality is even higher than that attained by the United States in 2000–2010, at least if capital gains are excluded; if they are included, the United States was slightly ahead of Colombia over the past decade.

It is important, however, to be aware of the significant limitations of the data available for measuring the dynamics of the income distribution in poor and emerging countries and for comparing them with the rich countries. The orders of magnitude indicated here are the best I was able to come up with given the available sources, but the truth is that our knowledge remains meager. We have tax data for the entire twentieth century for only a few emerging economies, and there are many gaps and breaks in the data, often in the period 1950–1970, the era of independence (in Indonesia, for example). Work is going forward to update the WTID with historical data from many other countries, especially from among the former British and French colonies, in Indochina and Africa, but data from the colonial era are often difficult to relate to contemporary tax records.

Where tax records do exist, their interest is often reduced by the fact that the income tax in less developed countries generally applies to only a small minority of the population, so that one can estimate the upper centile's share of total income but not the upper decile's. Where the data allow, as in South Africa for certain subperiods, one finds that the highest observed levels for the top decile are on the order of 50–55 percent of national income—a level comparable to or slightly higher than the

highest levels of inequality observed in the wealthy countries, in Europe in 1900–1910 and in the United States in 2000–2010.

I have also noticed a certain deterioration of the tax data after 1990. This is due in part to the arrival of computerized records, which in many cases led the tax authorities to interrupt the publication of detailed statistics, which in earlier periods they needed for their own purposes. This sometimes means, paradoxically, that sources have deteriorated since the advent of the information age (we find the same thing happening in the rich countries). Above all, the deterioration of the sources seems to be related to a certain disaffection with the progressive income tax in general on the part of certain governments and international organizations. A case in point is India, which ceased publishing detailed income tax data in the early 2000s, even though such data had been published without interruption since 1922. As a result, it is harder to study the evolution of top incomes in India since 2000 than over the course of the twentieth century.

This lack of information and democratic transparency is all the more regrettable in that the question of the distribution of wealth and of the fruits of growth is at least as urgent in the poor and emerging economies as in the rich ones. Note, too, that the very high official growth figures for developing countries (especially India and China) over the past few decades are based almost exclusively on production statistics. If one tries to measure income growth by using household survey data, it is often quite difficult to identify the reported rates of macroeconomic growth: Indian and Chinese incomes are certainly increasing rapidly, but not as rapidly as one would infer from official growth statistics. This paradox—sometimes referred to as the "black hole" of growth—is obviously problematic. It may be due to the overestimation of growth of output (there are many bureaucratic incentives for doing so), or perhaps the underestimation of income growth (household surveys have their own flaws), or most likely both. In particular, the missing income may be explained by the possibility that a disproportionate share of the growth in output has gone to the most highly remunerated individuals, whose incomes are not always captured in the tax data.

In the case of India, it is possible to estimate (using tax return data) that the increase in the upper centile's share of national income explains between one-quarter and one-third of the "black hole" of growth between 1990 and 2000. Given the deterioration of the tax data since 2000, it is impossible to do a proper social decomposition of recent growth. In the case of China, official tax records are even more rudimentary than in India. In the current state of research, the estimates in Figure 9.9 are the most reliable we have. It is nevertheless urgent that both countries publish more complete data —and other countries should do so as well. If and when better data become available, we may discover that inequality in India and China has increased more rapidly than we imagined.

In any case, the important point is that whatever flaws the tax authorities in poor and emerging countries may exhibit, the tax data reveal much higher—and more realistic—top income levels than do household surveys. For example, tax returns show that the top centile's share of national income in Colombia in 2000–2010 was more than 20 percent (and almost 20 percent in Argentina). Actual inequality may be even greater. But the fact that the highest incomes declared in household surveys in these same countries are generally only 4 to 5 times as high as the average income (suggesting that no one is really rich)—so that, if we were to trust the household survey, the top centile's share would be less than 5 percent—suggests that the survey data are not very credible. Clearly, household surveys,

which are often the only source used by international organizations (in particular the World Bank) and governments for gauging inequality, give a biased and misleadingly complacent view of the distribution of wealth. As long as these official estimates of inequality fail to combine survey data with other data systematically gleaned from tax records and other government sources, it will be impossible to apportion macroeconomic growth properly among various social groups or among the centiles and deciles of the income hierarchy. This is true, moreover, of wealthy countries as well as poor and emerging ones.

The Illusion of Marginal Productivity

Let me now return to the explosion of wage inequality in the United States (and to a lesser extent Britain and Canada) after 1970. As noted, the theory of marginal productivity and of the race between technology and education is not very convincing: the explosion of compensation has been highly concentrated in the top centile (or even the top thousandth) of the wage distribution and has affected some countries while sparing others (Japan and continental Europe are thus far much less affected than the United States), even though one would expect technological change to have altered the whole top end of the skill distribution in a more continuous way and to have worked its effects in all countries at a similar level of development. The fact that income inequality in the United States in 2000–2010 attained a level higher than that observed in the poor and emerging countries at various times in the past—for example, higher than in India or South Africa in 1920–1930, 1960–1970, and 2000–2010—also casts doubt on any explanation based solely on objective inequalities of productivity. Is it really the case that inequality of individual skills and productivities is greater in the United States today than in the half-illiterate India of the recent past (or even today) or in apartheid (or postapartheid) South Africa? If that were the case, it would be bad news for US educational institutions, which surely need to be improved and made more accessible but probably do not deserve such extravagant blame.

To my mind, the most convincing explanation for the explosion of the very top US incomes is the following. As noted, the vast majority of top earners are senior managers of large firms. It is rather naïve to seek an objective basis for their high salaries in individual "productivity." When a job is replicable, as in the case of an assembly-line worker or fast-food server, we can give an approximate estimate of the "marginal product" that would be realized by adding one additional worker or waiter (albeit with a considerable margin of error in our estimate). But when an individual's job functions are unique, or nearly so, then the margin of error is much greater. Indeed, once we introduce the hypothesis of imperfect information into standard economic models (eminently justifiable in this context), the very notion of "individual marginal productivity" becomes hard to define. In fact, it becomes something close to a pure ideological construct on the basis of which a justification for higher status can be elaborated.

To put this discussion in more concrete terms, imagine a large multinational corporation employing 100,000 people and with gross annual revenue of 10 billion euros, or 100,000 euros per worker. Suppose that half of this revenue figure represents purchases of goods and services by the firm (this is a typical figure for the economy as a whole), so that the value added by the firm—the value available to pay the labor and capital that it directly employs—is 5 billion euros, or 50,000 euros per worker.

To set the pay of the firm's CFO (or his deputies, or of the director of marketing and her staff, etc.), one would in principle want to estimate his marginal productivity, that is, his contribution to the firm's value-added of 5 billion euros: is it 100,000, 500,000, or 5 million euros per year? A precise, objective answer to this question is clearly impossible. To be sure, one could in theory experiment by trying out several CFOs, each for several years, in order to determine what impact the choice has on the firm's total revenue of 10 billion euros. Obviously, such an estimate would be highly approximate, with a margin of error much greater than the maximum salary one would think of paying, even in a totally stable economic environment. And the whole idea of experimentation looks even more hopeless when one remembers that the environment is in fact changing constantly, as is the nature of the firm and the exact definition of each job.

In view of these informational and cognitive difficulties, how are such remunerations determined in practice? They are generally set by hierarchical superiors, and at the very highest levels salaries are set by the executives themselves or by corporate compensation committees whose members usually earn comparable salaries (such as senior executives of other large corporations). In some companies, stockholders are asked to vote on compensation for senior executives at annual meetings, but the number of posts subject to such approval is small, and not all senior managers are covered. Since it is impossible to give a precise estimate of each manager's contribution to the firm's output, it is inevitable that this process yields decisions that are largely arbitrary and dependent on hierarchical relationships and on the relative bargaining power of the individuals involved. It is only reasonable to assume that people in a position to set their own salaries have a natural incentive to treat themselves generously, or at the very least to be rather optimistic in gauging their marginal productivity. To behave in this way is only human, especially since the necessary information is, in objective terms, highly imperfect. It may be excessive to accuse senior executives of having their "hands in the till," but the metaphor is probably more apt than Adam Smith's metaphor of the market's "invisible hand." In practice, the invisible hand does not exist, any more than "pure and perfect" competition does, and the market is always embodied in specific institutions such as corporate hierarchies and compensation committees.

This does not mean that senior executives and compensation committees can set whatever salaries they please and always choose the highest possible figure. "Corporate governance" is subject to certain institutions and rules specific to each country. The rules are generally ambiguous and flawed, but there are certain checks and balances. Each society also imposes certain social norms, which affect the views of senior managers and stockholders (or their proxies, who are often institutional investors such as financial corporations and pension funds) as well as of the larger society. These social norms reflect beliefs about the contributions that different individuals make to the firm's output and to economic growth in general. Since uncertainty about these issues is great, it is hardly surprising that perceptions vary from country to country and period to period and are influenced by each country's specific history. The important point is that it is very difficult for any individual firm to go against the prevailing social norms of the country in which it operates.

Without a theory of this kind, it seems to me quite difficult to explain the very large differences of executive pay that we observe between on the one hand the United States (and to a lesser extent in other English-speaking countries) and on the other continental Europe and Japan. Simply put, wage

inequalities increased rapidly in the United States and Britain because US and British corporations became much more tolerant of extremely generous pay packages after 1970. Social norms evolved in a similar direction in European and Japanese firms, but the change came later (in the 1980s or 1990s) and has thus far not gone as far as in the United States. Executive compensation of several million euros a year is still more shocking today in Sweden, Germany, France, Japan, and Italy than in the United States or Britain. It has not always been this way—far from it: recall that in the 1950s and 1960s the United States was more egalitarian than France, especially in regard to the wage hierarchy. But it has been this way since 1980, and all signs are that this change in senior management compensation has played a key role in the evolution of wage inequalities around the world.

The Takeoff of the Supermanagers: A Powerful Force for Divergence

This approach to executive compensation in terms of social norms and acceptability seems rather plausible a priori, but in fact it only shifts the difficulty to another level. The problem is now to explain where these social norms come from and how they evolve, which is obviously a question for sociology, psychology, cultural and political history, and the study of beliefs and perceptions at least as much as for economics per se. The problem of inequality is a problem for the social sciences in general, not for just one of its disciplines. In the case in point, I noted earlier that the "conservative revolution" that gripped the United States and Great Britain in the 1970s and 1980s, and that led to, among other things, greater tolerance of very high executive pay, was probably due in part to a feeling that these countries were being overtaken by others (even though the postwar period of high growth in Europe and Japan was in reality an almost mechanical consequence of the shocks of the period 1914–1945). Obviously, however, other factors also played an important role.

To be clear, I am not claiming that all wage inequality is determined by social norms of fair remuneration. As noted, the theory of marginal productivity and of the race between technology and education offers a plausible explanation of the long-run evolution of the wage distribution, at least up to a certain level of pay and within a certain degree of precision. Technology and skills set limits within which most wages must be fixed. But to the extent that certain job functions, especially in the upper management of large firms, become more difficult to replicate, the margin of error in estimating the productivity of any given job becomes larger. The explanatory power of the skills-technology logic then diminishes, and that of social norms increases. Only a small minority of employees are affected, a few percent at most and probably less than 1 percent, depending on the country and period.

But the key fact, which was by no means evident a priori, is that the top centile's share of total wages can vary considerably by country and period, as the disparate evolutions in the wealthy countries after 1980 demonstrate. The explosion of supermanager salaries should of course be seen in relation to firm size and to the growing diversity of functions within the firm. But the objectively complex problem of governance of large organizations is not the only issue. It is also possible that the explosion of top incomes can be explained as a form of "meritocratic extremism," by which I mean the apparent need of modern societies, and especially US society, to designate certain individuals as "winners" and to reward them all the more generously if they seem to have been selected on the basis of their intrinsic merits rather than birth or background. (I will come back to this point.)

In any case, the extremely generous rewards meted out to top managers can be a powerful force for

divergence of the wealth distribution: if the best paid individuals set their own salaries, (at least to some extent), the result may be greater and greater inequality. It is very difficult to say in advance where such a process might end. Consider again the case of the CFO of a large firm with gross revenue of 10 billion euros a year. It is hard to imagine that the corporate compensation committee would suddenly decide that the CFO's marginal productivity is 1 billion or even 100 million euros (if only because it would then be difficult to find enough money to pay the rest of the management team). By contrast, some people might think that a pay package of 1 million, 10 million, or even 50 million euros a year would be justified (uncertainty about individual marginal productivity being so large that no obvious limit is apparent). It is perfectly possible to imagine that the top centile's share of total wages could reach 15–20 percent in the United States, or 25–30 percent, or even higher.

The most convincing proof of the failure of corporate governance and of the absence of a rational productivity justification for extremely high executive pay is that when we collect data about individual firms (which we can do for publicly owned corporations in all the rich countries), it is very difficult to explain the observed variations in terms of firm performance. If we look at various performance indicators, such as sales growth, profits, and so on, we can break down the observed variance as a sum of other variances: variance due to causes external to the firm (such as the general state of the economy, raw material price shocks, variations in the exchange rate, average performance of other firms in the same sector, etc.) plus other "nonexternal" variances. Only the latter can be significantly affected by the decisions of the firm's managers. If executive pay were determined by marginal productivity, one would expect its variance to have little to do with external variances and to depend solely or primarily on nonexternal variances. In fact, we observe just the opposite: it is when sales and profits increase for external reasons that executive pay rises most rapidly. This is particularly clear in the case of US corporations: Bertrand and Mullainhatan refer to this phenomenon as "pay for luck."

I return to this question and generalize this approach in Part Four (see Chapter 14). The propensity to "pay for luck" varies widely with country and period, and notably as a function of changes in tax laws, especially the top marginal income tax rate, which seems to serve either as a protective barrier (when it is high) or an incentive to mischief (when it is low)—at least up to a certain point. Of course changes in tax laws are themselves linked to changes in social norms pertaining to inequality, but once set in motion they proceed according to a logic of their own. Specifically, the very large decrease in the top marginal income tax rate in the English-speaking countries after 1980 (despite the fact that Britain and the United States had pioneered nearly confiscatory taxes on incomes deemed to be indecent in earlier decades) seems to have totally transformed the way top executive pay is set, since top executives now had much stronger incentives than in the past to seek large raises. I also analyze the way this amplifying mechanism can give rise to another force for divergence that is more political in nature: the decrease in the top marginal income tax rate led to an explosion of very high incomes, which then increased the political influence of the beneficiaries of the change in the tax laws, who had an interest in keeping top tax rates low or even decreasing them further and who could use their windfall to finance political parties, pressure groups, and think tanks.