

Putting Distribution Back at the Center of Economics: Reflections on *Capital in the Twenty-First Century*

Thomas Piketty

When a lengthy book is widely discussed in academic circles and the popular media, it is probably inevitable that the arguments of the book will be simplified in the telling and retelling. In the case of my book *Capital in the Twenty-First Century* (2014), a common simplification of the main theme is that because the rate of return on capital r exceeds the growth rate of the economy g , the inequality of wealth is destined to increase indefinitely over time. In my view, the magnitude of the gap between r and g is indeed one of the important forces that can explain historical magnitudes and variations in wealth inequality: in particular, it can explain why wealth inequality was so extreme and persistent in pretty much every society up until World War I (for discussion, see Chapter 10 of my book). That said, the way in which I perceive the relationship between $r > g$ and wealth inequality is often not well-captured in the discussion that has surrounded my book—even in discussions by research economists.

In this essay, I will return to some of the themes of my book and seek to clarify and refocus the discussion concerning those themes. For example, I do not view $r > g$ as the only or even the primary tool for considering changes in income and wealth in the 20th century, or for forecasting the path of income and wealth inequality in the 21st century. Institutional changes and political shocks—which can be viewed as largely endogenous to the inequality and development process itself—played a major role in the past, and will probably continue to do so in the future. In addition, I certainly do not believe that $r > g$ is a useful tool for the discussion of rising inequality of labor income: other mechanisms

■ Thomas Piketty is Professor of Economics at the Paris School of Economics, Paris, France. His email address is piketty@psemail.eu.

and policies are much more relevant here, for example, the supply and demand of skills and education. One of my main conclusions is that there is substantial uncertainty about how far income and wealth inequality might rise in the 21st century and that we need more transparency and better information about income and wealth dynamics so that we can adapt our policies and institutions to a changing environment.

My book is primarily about the history of the distribution of income and wealth. Thanks to the cumulative efforts of several dozen scholars, we have been able to collect a relatively large historical database on the structure of national income and national wealth, and the evolution of income and wealth distributions, covering three centuries and over 20 countries. The first objective of my book was to present this body of historical evidence and to analyze the economic, social, and political processes that can account for the evolutions that we observe in the various countries since the Industrial Revolution. I stress from the beginning that we have too little historical data at our disposal to be able to draw definitive judgments. On the other hand, at least we have substantially more evidence than we used to have.

My book is probably best described as an analytical historical narrative based upon this new body of evidence. In this way, I hope I can contribute to placing the study of distribution and of the long-run back at the center of economic thinking. Many 19th century economists, including Thomas Malthus, David Ricardo, and Karl Marx, put the distribution question at the center of political economy. However, they had limited data at their disposal, and so their approach was mostly theoretical. In contrast, since the mid-20th century, a number of economists, most notably Simon Kuznets and Anthony Atkinson, have been developing the possibility of an approach that blends theory with more data-intensive and historical approaches. This historical data collection project on which my book is based follows directly in the tradition of the pioneering works by Kuznets (1953) and Atkinson and Harrison (1978).

In this essay, I will take up several themes from my book that have perhaps become attenuated or garbled in the ongoing discussions of the book, and will seek to re-explain and re-frame these themes. First, I stress the key role played in my book by the interaction between beliefs systems, institutions, and the dynamics of inequality. Second, I briefly describe my multidimensional approach to the history of capital and inequality. Third, I review the relationship and differing causes between wealth inequality and income inequality. Fourth, I turn to the specific role of $r > g$ in the dynamics of wealth inequality: specifically, a larger $r - g$ gap will amplify the steady-state inequality of a wealth distribution that arises out of a given mixture of shocks. Fifth, I consider some of the scenarios that affect how $r - g$ might evolve in the 21st century, including rising international tax competition, a growth slowdown, and differential access by the wealthy to higher returns on capital. Finally, I seek to clarify what is distinctive in my historical and political economy approach to institutions and inequality dynamics, and the complementarity with other approaches.

Beliefs Systems, Institutions, and the Dynamics of Inequality

In my book, I attempt to study not only the dynamics of income and wealth inequality, but also the evolution of collective representations of social inequality in public discussions and political debates, as well as in literature and movies. I believe that the analysis of representations and beliefs systems about income and wealth is an integral and indispensable part of the study of income and wealth dynamics.

Indeed, a main conclusion of my analytical historical narrative is stated in the introduction of the book (p. 20, 35), that “one should be wary of any economic determinism in regard to inequalities of wealth and income . . . The history of the distribution of wealth has always been deeply political, and it cannot be reduced to purely economic mechanisms. . . . It is shaped by the way economic, social, and political actors view what is just and what is not, as well as by the relative power of those actors and the collective choices that result. It is the joint product of all relevant actors combined. . . . How this history plays out depends on how societies view inequalities and what kinds of policies and institutions they adopt to measure and transform them.” As I wrote in a follow-up essay with a co-author: “In a sense, both Marx and Kuznets were wrong. There are powerful forces pushing alternatively in the direction of rising or shrinking inequality. Which one dominates depends on the institutions and policies that societies choose to adopt” (Piketty and Saez 2014, p. 842–43).

The role of political shocks and changing representations of the economy is especially obvious when one studies inequality dynamics during the 20th century. In particular (p. 20), “the reduction of inequality that took place in most developed countries between 1910 and 1950 was above all a consequence of war and revolution and of policies adopted to cope with these shocks. Similarly, the resurgence of inequality after 1980 is due largely to the opposite political shifts of the past several decades, especially in regard to taxation and finance.”

I also try to show that belief systems about the distribution of income and wealth matter a great deal if one wants to understand the structure of inequality in the 18th and 19th centuries, and indeed in any society. Each country has its own intimate history with inequality, and I attempt to show that national identities play an important role in the two-way interaction between inequality dynamics and the evolution of perceptions, institutions, and policies.

I continually refer to a large number of other institutions and public policies that play a substantial role in my historical account of inequality dynamics across three centuries and over 20 countries. I emphasize the importance of educational institutions (in particular the extent of equal access to high-quality schools and universities) and of fiscal institutions (especially the chaotic advent of progressive taxation of income, inheritance, and wealth). Other examples of important factors include: the development of the modern welfare state; monetary regimes, central banking, and inflation; labor market rules, minimum wages, and collective bargaining; forced labor (slavery); colonialism, wars, and revolutions; expropriations, physical destruction, and privatizations; corporate governance and

stakeholder rights; rent and other price controls (such as the prohibition or limitation of usury); financial deregulation and capital flows; trade policies; family transmission rules and legal property regimes; fertility policies; and many others.

A Multidimensional History of Capital and Inequality

A central reason that my book is relatively long is that I try to offer a relatively detailed, multidimensional history of capital and its metamorphosis. Capital ownership takes many different historical forms, and each of them involves different forms of institutions, rules, and power relations, which must be analyzed as such.

Theoretical models, abstract concepts, and equations (such as $r > g$, to which I return in greater detail below) also play a certain role in my analysis. However this role is relatively modest—as I believe the role of theory should generally be in the social sciences—and it should certainly not be exaggerated. Models can contribute to clarifying logical relationships between particular assumptions and conclusions but only by oversimplifying the real world to an extreme point. Models can play a useful role but only if one does not overestimate the meaning of this kind of abstract operation. All economic concepts, irrespective of how “scientific” they pretend to be, are intellectual constructions that are socially and historically determined, and which are often used to promote certain views, values, or interests. Models are a language that can be useful only if solicited together with other forms of expressions, while recognizing that we are all part of the same conflict-filled, deliberative process.

In particular, the notion of an aggregate capital stock K and of an aggregate production function $Y = F(K, L)$ are highly abstract concepts. From time to time, I refer to them. But I certainly do not believe that such grossly oversimplified concepts can provide an adequate description of the production structure and the state of property and social relations for any society. For example, I explain in Chapter 1, when I define capital and wealth (p. 47):

Capital is not an immutable concept: it reflects the state of development and prevailing social relations of each society. . . . The boundary between what private individuals can and cannot own has evolved considerably over time and around the world, as the extreme case of slavery indicates. The same is true of property in the atmosphere, the sea, mountains, historical monuments, and knowledge. Certain private interests would like to own these things, and sometimes they justify this desire on grounds of efficiency rather than mere self-interest. But there is no guarantee that this desire coincides with the general interest.

More generally, I analyze the diversity of the forms taken by capital assets and the problems raised by property relations and market valorizations throughout

history. I study in some length the many transformations in the nature of capital assets, from agricultural land to modern real estate and business and financial capital. Each type of asset has its own particular economic and political history and gives rise to different bargaining processes, power struggles, economic innovations, and social compromises.

For example, the fact that capital ownership and property rights are historically determined is particularly clear when I study the role of slave capital in the Southern United States before 1865, which can be viewed as the most extreme form of ownership and domination of owners over others (Chapter 4). A similar theme also becomes evident when I examine the lower stock market capitalization of German companies relative to their Anglo-American counterparts, a phenomenon that is certainly related to the fact that German shareholders need to share power with other stakeholders (workers, governments, nongovernment organizations, and others) somewhat more than in other countries (Chapter 5). This power-sharing apparently is not detrimental to the productive efficiency and exporting performance of German firms, which illustrates the fact that the market and social values of capital can often differ.

Other examples involve real estate capital and natural resource wealth—like oil. Large upward or downward movements of real estate prices play an important role in the evolution of aggregate capital values during recent decades, as they did during the first half of the 20th centuries (in particular, Chapters 3–6). This can in turn be accounted for by a complex mixture of institutional and technological forces, including rent control policies and other rules regulating relations between owners and tenants, the transformation of economic geography, and the changing speed of technical progress in the transportation and construction industries relative to other sectors. The issue of oil capital and its world distribution is rooted in the power relations and military protections that go with it (in particular in the Middle East), which also have consequences for the financial investment strategies followed by the corresponding sovereign wealth funds (discussed in Chapter 12).

The institutional analysis of property relations and capital assets also has international and public-sector dimensions. The hypertrophy of gross financial asset positions between countries, which is one of the main characteristics of the financial globalization process of recent decades, is a recurring theme of the book (Chapters 1–5, 12, 15, and 16). I analyze the very large magnitude of the net foreign assets positions reached by Britain and France at the height of their colonial empires, and I compare them to today's net positions of China, Japan, or Germany. I repeatedly stress that international property relations—the fact that economic actors in some countries own significant claims on real and financial assets in other countries—can be particularly complicated to regulate in a peaceful manner. This was certainly true during the colonization and decolonization periods. Issues of international property relations could erupt again in the future. The difficulty in dealing with extreme internal and external inequality certainly contributes to explaining the high political instability that has long plagued the development process in Latin American and African countries.

Public capital—which depends on the changing patterns and complex political histories of public investment and deficit trajectories and nationalization and privatization policies—also plays a critical role in the book (especially Chapters 3 and 4). I emphasize the sharp dissimilarities in country experiences (contrasting in particular the cases of Britain and France in the 18th and 19th centuries), as well as the commonalities (such as the historically large level of public capital in the post-World War II period, and the large decline in recent decades in high-income countries as well as in Russia or China, with important consequences for the distribution of private wealth and the rise of new forms of oligarchs).

Given the specific and context-heavy discussion of these multidimensional factors, does it still make sense to speak of “capital” as a single category? The fact that it is technically possible to add up all the market values of the different existing assets (to the extent that such market values are well defined, which is not always entirely clear) in order to compute the aggregate value of the capital stock K does not change anything about the basic multidimensional reality of assets and corresponding property relations. I attempt to show that this abstract operation can be useful for some purposes. In particular, by computing the ratio $\beta = K/Y$ between the aggregate market value of capital K and national income Y , one can compare the overall importance of capital wealth, private property, and public property in societies that are otherwise impossible to compare. For instance, one finds that in spite of all metamorphosis in the nature of assets and institutional arrangements, aggregate capital values—expressed relative to total national income—are in a number of countries approaching the levels observed in the patrimonial societies that flourished in the 18th–19th centuries and until World War I. I believe that this finding is interesting in itself. But it certainly does not alter the fact that a proper comparison of these different societies requires a careful separate analysis of the various asset categories and corresponding social and economic relations.

Inequality of Labor Income and Inequality of Wealth

Another way in which my analysis of capital and inequality is multidimensional is that throughout the book, I continually distinguish between the inequality of labor income and the inequality of capital ownership. Of course these two dimensions of inequality do interact in important ways: for example, rising inequality in labor earnings during a certain period of time might tend to fuel rising wealth concentration in following decades or generations. But the forces that drive income inequality and wealth inequality are largely different.

For instance, I point out in my book (particularly Chapters 8–9) that the rise of top income shares in the United States over the 1980–2010 period is due for the most part to rising inequality of labor earnings, which can itself be explained by a mixture of two groups of factors: 1) rising inequality in access to skills and to higher education over this time period in the United States, an evolution which might

have been exacerbated by rising tuition fees and insufficient public investment; and 2) exploding top managerial compensation, itself probably stimulated by changing incentives and norms, and by large cuts in top tax rates (see also Chapter 14; Piketty, Saez, and Stantcheva 2014). More broadly, I argue (p. 243) that the mechanisms behind unequal incomes from labor “include the supply of and demand for different skills, the state of the educational system, and the various rules and institutions that affect the operation of the labor market and the determination of wages.” This rise in labor earnings inequality in recent decades evidently has little to do with the gap $r - g$; indeed, it seems fairly difficult to find a logical way that $r - g$ could affect the inequality of labor income. Conversely, “[i]n the case of unequal incomes from capital, the most important processes involve savings and investment behavior, laws governing gift-giving and inheritance, the operation of real estate and financial markets, and so on” (p. 243).

In addition, the notions of top deciles or percentiles are not the same for the distributions of labor income and capital ownership. The use of deciles and percentages should be viewed as a language allowing for comparisons between societies that are otherwise impossible to compare, such as France in 1789 and China or the United States in 2014, in the same way as the aggregate capital-income ratio can be used to make comparisons. But in certain societies, the top shares of income and wealth might be highly correlated, while in other societies they may represent entirely different social hierarchies (as in traditional patriarchal societies). The extent to which these two dimensions of inequality differ gives rise to different representations and beliefs systems about social inequality, which in turn shape institutions and public policies affecting inequality dynamics.

The Dynamics of Wealth Inequality and the Role of $r > g$

Let me now try to clarify the role played by $r > g$ in my analysis of inequality dynamics. The rate of return on capital is given by r , while g measures the rate of economic growth. The gap between r and g is certainly not the only relevant mechanism for analyzing the dynamics of wealth inequality. As I explained in the previous sections, a wide array of institutional factors are central to understanding the evolution of wealth.

Moreover, the insight that the rate of return to capital r is permanently higher than the economy’s growth rate g does not in itself imply anything about wealth inequality. Indeed the inequality $r > g$ holds true in the steady-state equilibrium of most standard economic models, including in representative-agent models where each individual owns an equal share of the capital stock.

For instance, consider the standard dynamic model where each individual behaves as an infinitely lived family and where the steady-state rate of return is well known to be given by the modified “golden rule” $r = \theta + \gamma g$ (where θ is the rate of time preference and γ is the curvature of the utility function). For example, if

$\theta = 3$ percent, $\gamma = 2$, and $g = 1$ percent, then $r = 5$ percent.¹ In this framework, the inequality $r > g$ always holds true, and this does not entail any implication about wealth inequality.

In a representative agent framework, what $r > g$ means is that in steady-state each family only needs to reinvest a fraction g/r of its capital income in order to ensure that its capital stock will grow at the same rate g as the size of the economy, and the family can then consume a fraction $1 - g/r$. For example, if $r = 5$ percent and $g = 1$ percent, then each family will reinvest 20 percent of its capital income and can consume 80 percent. Again, $r > g$, but this tells us nothing at all about inequality: this is simply saying that capital ownership allows the economy to reach higher consumption levels—which is really the very least one can ask from capital ownership.²

So what is the relationship between $r - g$ and wealth inequality? To answer this question, one needs to introduce extra ingredients into the basic model so that inequality arises in the first place.³ In the real world, many shocks to the wealth trajectories of families can contribute to making the wealth distribution highly unequal (indeed, in every country and time period for which we have data, wealth distribution *within each age group* is substantially more unequal than income distribution, which is difficult to explain with standard life-cycle models of wealth accumulation; for a concise summary of the historical evidence on the extent of income and wealth inequality, see Piketty and Saez 2014). There are demographic shocks: some families have many children and have to split inheritances in many pieces, some have few; some parents die late, some die soon; and so on. There are also shocks to rates of return: some families make very good investments, others go bankrupt. There are shocks to labor market outcomes: some earn high wages, others do not. There are differences in taste parameters that affect the level of saving: some families consume

¹ Intuitively, in a model where everyone maximizes an infinite-horizon utility function $U = \int_{0 \leq t \leq +\infty} e^{-\theta t} u(c_t)$ (with $u(c) = c^{1-\gamma}/(1-\gamma)$), then $r = \theta + \gamma g$ is the unique rate of return to capital possible in the long-run for the following reason: it is the sole rate such that the agents are willing to raise their consumption at rate g , that is at the growth rate of the economy. If the return is higher, the agents prefer to postpone their consumption and accumulate more capital, which will decrease the rate of return; and if it is lower, they want to anticipate their consumption and borrow more, which will increase the rate of return.

² The inequality $r < g$ would correspond to a situation which economists often refer to as “dynamic inefficiency”: in effect, one would need to invest more than the return to capital in order to ensure that one’s capital stock keeps rising as fast as the size of the economy. In infinite horizon models with perfect capital markets, this cannot happen. In effect, $r < g$ would violate the transversality condition: the net present value of future resources would be infinite, so that rational agents would borrow infinite amounts in order to consume right away. However, in models with other saving motives, such as finite-horizon overlapping generation models, it is possible for $r < g$.

³ In the dynamic model with no shock, there is no force generating inequality out of equality (or equality out of inequality), so that any initial level of wealth inequality (including full equality) can be self-sustaining, as long as the modified “golden rule” is satisfied. In effect, steady-state wealth inequality is exogenous and indeterminate, and does not depend on the gap $r - g$. Note however that the magnitude of the gap $r - g$ has an effect on the steady-state inequality of consumption and welfare in this basic model: for example, if $r - g$ is small, then high-wealth dynasties need to reinvest a large fraction of their capital income, so that they do not consume much more than low-wealth dynasties.

a lot more than a fraction $1 - g/r$ of their capital income and might even consume away the capital value and die with negligible wealth; others might reinvest a lot more than a fraction g/r and have a strong taste for leaving bequests and perpetuating large fortunes.

A central property of this large class of models is that for a given structure of shocks, the long-run magnitude of wealth inequality will tend to be magnified if the gap $r - g$ is higher. In other words, wealth inequality will converge towards a finite level in these models. The shocks will ensure that there is always some degree of downward and upward wealth mobility such that wealth inequality remains bounded in the long run. But this finite inequality level will be a steeply rising function of the gap $r - g$. Intuitively, a higher gap between r and g works as an amplifier mechanism for wealth inequality for a given variance of other shocks. To put it differently: a higher gap between r and g allows an economy to sustain a level of wealth inequality that is higher and more persistent over time (that is, a higher gap $r - g$ leads both to higher inequality and lower mobility).

More precisely, one can show that if shocks take a multiplicative form, then in the long run, the inequality of wealth will converge toward a distribution that has a Pareto shape for top wealth holders (which is approximately the form that we observe in real-world distributions and corresponds to relatively fat upper tails and a large concentration of wealth at the very top), and that the inverted Pareto coefficient (an indicator of top-end inequality) is a steeply rising function of the gap $r - g$.⁴ This well-known theoretical result was established by a number of authors using various structures of demographic and economic shocks (see in particular Champernowne 1953; Stiglitz 1969). The logic behind this result and this “inequality amplification” impact of $r - g$ is presented in Chapter 10 of my book: for detailed references to this literature on wealth inequality, $r - g$, and Pareto coefficients see the online appendix to Chapter 10 of my book (available at <http://piketty.pse.ens.fr/capital21c>) and Piketty and Zucman (2015, section 5.4). These connections between $r - g$ and Pareto coefficients of steady-state wealth distributions are also explained very clearly in the review by Charles Jones in the present symposium.

In this class of models, relatively small changes in $r - g$ can generate very large changes in steady-state wealth inequality. For example, simple simulations of the model with binomial taste shocks show that going from $r - g = 2$ percent to $r - g = 3$ percent is sufficient to move the inverted Pareto coefficient from $b = 2.28$ to $b = 3.25$. This corresponds to a shift from an economy with moderate wealth inequality—say, with a top 1 percent wealth share around 20–30 percent, such as present-day Europe or the United States—to an economy with very high wealth

⁴ A Pareto distribution means that above a certain wealth level z_0 , the population fraction with wealth above z is given by $p(z) = p_0(z_0/z)^a$ (where a is a constant). A characteristic property of the Pareto distribution is that the ratio $b = E(z|z > z')/z'$ between average wealth above some threshold z' and the level of the threshold z' is independent of z' and is equal to the inverted Pareto coefficient $b = a/(a - 1)$.

inequality, with a top 1 percent wealth share around 50–60 percent, such as Europe in the 18th–19th centuries and up until World War I.⁵

To summarize: the effect of $r - g$ on inequality follows from its dynamic cumulative effects in wealth accumulation models with random shocks, and the quantitative magnitude of this impact seems to be sufficiently large to account for very important variations in wealth inequality.

To reiterate, this argument does not imply that the $r - g$ effect is the only important force that matters in accounting for historical variations in wealth inequality. The variance of other shocks (particularly to rates of returns, which vary enormously across assets and individuals), as well as the income and wealth profiles of saving rates, obviously matter a great deal. Most importantly, it is really the *interaction* between the $r - g$ effect and the institutional and public policy responses—including progressive taxation of income, wealth, and inheritance; inflation; nationalizations, physical destruction, and expropriations; estate division rules; and so on—which in my view, determines the dynamics and the magnitude of wealth inequality. In particular, if one introduces taxation into the basic model, then it follows immediately that what determines long-run wealth inequality and the steady-state Pareto coefficient is the gap $(1 - t)r - g$ between the net-of-tax rate of return and the growth rate.

In their contribution to this symposium, Acemoglu and Robinson present cross-country regression results between income inequality and $r - g$ and argue that $r - g$ does not seem to have much impact on inequality. However, I do not find these regressions very convincing, for two main reasons. First, income inequality is primarily determined by the inequality of labor income (which typically represents between two-thirds and three-quarters of total income), which as I noted above has nothing to do with $r - g$. It would make more sense to run such a regression with wealth inequality, but long-run wealth inequality series are available for a much more limited number of countries than income inequality series. In Chapter 12 of my book, I present wealth inequality series for only four countries (France, Britain, Sweden, and the United States), and the data are far from perfect. We do plan in the future to extend the World Top Incomes Database (WTID) into a World Wealth and Income Database (W2ID) and to provide homogenous wealth inequality series for all countries covered in the WTID (over 30 countries). But at this stage, we have to do with what we have.

⁵ In the special case with saving taste shocks, the transition equation for normalized wealth $z_{ti} = w_{ti}/w_t$ (where w_{ti} is the wealth level of dynasty i at period t , and w_t is average wealth at period t) is given by: $z_{t+1i} = (s_{ti}/s) \cdot [(1 - \omega) + \omega \cdot z_{ti}]$, with $\omega = s \cdot e^{(r-g)H}$ (where s is the average saving taste parameter, s_{ti} is the taste parameter of dynasty i at period t , r and g are the annual rate of return and growth rate, and H is generation length). With binomial shocks with probability p , one can show that the inverted Pareto coefficient is given by $b = \log(1/p)/\log(1/\omega)$. See Piketty and Zucman (2015, section 5.4) for calibrations of this formula. In Atkinson, Piketty, and Saez (2011, figures 12–15, p. 50–55), we provide evidence on the long-run evolution of inverted Pareto coefficients for income distributions. See also the discussion in the online appendix to Chapter 10 of my book (available at <http://piketty.pse.ens.fr/capital21c>).

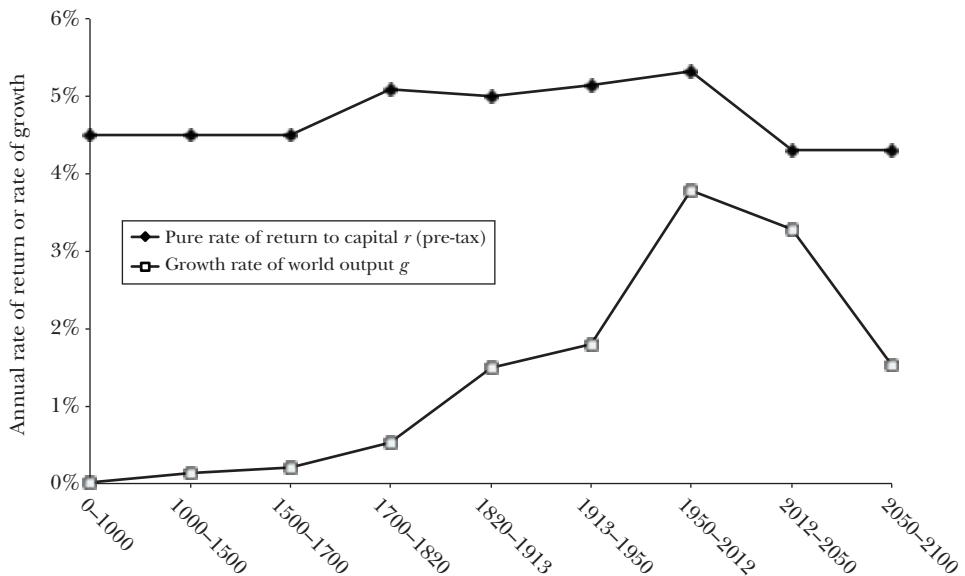
Second, the process of intergenerational accumulation and distribution of wealth is very long-run process, so looking at cross-sectional regressions between inequality and $r - g$ may not be very meaningful. One would need to introduce time lags, possibly over very long time periods: for example, one might use the average $r - g$ observed over 30 or 50 years. As I argue below, the broad correlations between $r - g$ and wealth inequality certainly seem to run in the right direction, both from a long-run (18th–19th versus 20th centuries) and international (Europe versus US) perspective. However, given the data limitations and the time-lag specification problems, I am not sure there is a lot to learn from running explicit cross-country regressions.

In my view, a more promising approach—on this issue as well as on many other issues—is a mixture of careful case studies and structural calibrations of theoretical models. Although we do not have many historical series on wealth inequality, they show a consistent pattern. Namely, we observe extremely high concentration of wealth in pretty much every European society in the 18th and 19th centuries up until World War I. In particular, in France, Britain, and Sweden, the top 10 percent wealth share was about 90 percent of total wealth (including a top 1 percent wealth share of around 60–70 percent) in the 19th century and at the very beginning of the 20th century. If anything, wealth inequality seems to have been rising somewhat during the 19th century and up until World War I—or maybe to have stabilized at very high levels around 1890–1910. Thus, in spite of the large changes in the nature of wealth during the 19th century—agricultural land as a form of wealth is largely replaced by real estate, business assets, and foreign investment—wealth inequality was as extreme in the modern industrial society of 1914 as it had been under France’s *ancien régime* in 1789.

The most convincing explanation for the very high wealth concentration in these pre–World War I European societies seems to be the very large $r - g$ gap—that is, the gap between rates of return and growth rates during the 18th and 19th centuries. There was very little taxation or inflation up until 1914, so the gap $(1 - t)r - g$ was particularly high in pre–World War I societies, which in dynamic models of wealth accumulation with random shocks leads to very large wealth concentration. In contrast, following the large capital shocks of the 1914–1945 period—a time of physical destruction, periods of high inflation and taxation, and nationalizations—the after-tax, after-capital-losses rate of return precipitously fell below the growth rate after World War I. Figure 1 compares the pre-tax pure rate of return with growth rate g , while Figure 2 shows a post-tax, post-losses rate of return, including projections into the future.

This interpretation of the evidence is further confirmed by the detailed individual-level data collected in French inheritance archives since the time of the French Revolution (Piketty, Postel-Vinay, and Rosenthal 2006, 2014). We find that the more and more steeply increasing age-wealth profiles at high wealth levels in the 19th century and early 20th century can be well accounted for by a capitalization effect and a high gap between $(1 - t)r$ and g . This age–wealth pattern suddenly breaks down following the 1914–1945 capital shocks. The fact that US

Figure 1

Rate of Return versus Growth Rate at the World Level, from Antiquity until 2100

Source: Author (figure 10.9 from Piketty 2014). For more on sources and series, see <http://piketty.pse.ens.fr/capital21c>.

Note: The rate of return to capital (pre-tax) has always been higher than the world growth rate, but the gap was reduced during the 20th century, and might widen again in the 21st century.

wealth concentration was significantly less than in Europe during the 19th century and up until World War I is also consistent with this model: growth rates were higher in the US economy, in particular due to higher population growth, thereby limiting the dynamic cumulative effects of the inequality amplification channel. Also, there had been less time for dynastic wealth concentration to arise in the US economy by the 19th century. This evidence is further reviewed in Chapters 10–11 of my book.

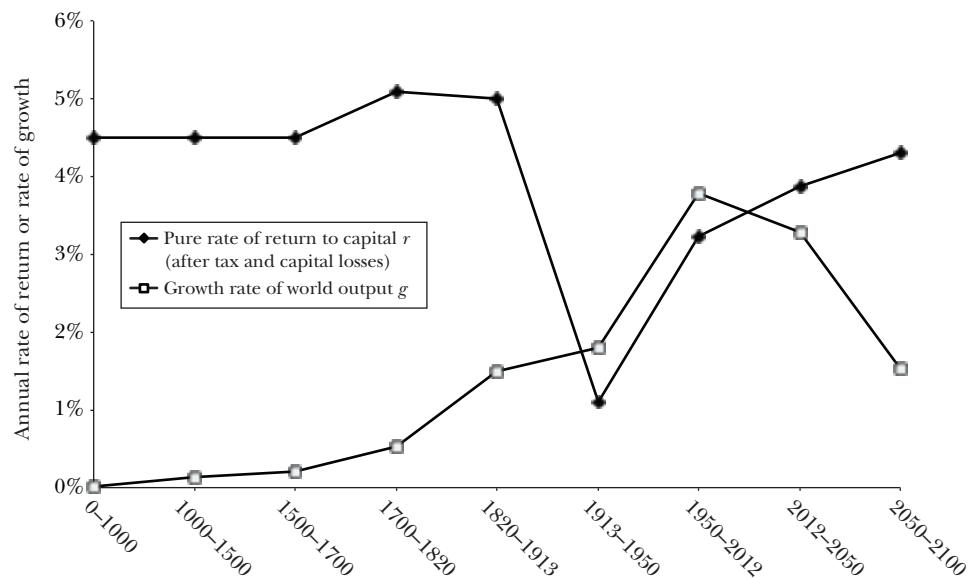
Data collection in French archives and in other countries will continue, and new data will certainly allow for better empirical tests of wealth accumulation models in the future. But at this stage, the best evidence we have suggests that $r > g$ is an important part of the explanation for the very high and persistent level of wealth concentration that we observe in most societies in the 18th–19th centuries and up until World War I.

What Will Be the Evolution of $r - g$ in the 21st Century?

A number of forces might lead to greater inequality of wealth in the 21st century, including a rise in the variance of shocks to demographic factors, rates of return,

Figure 2

After-Tax Rate of Return versus Growth Rate at the World Level, from Antiquity until 2100



Source: Author (figure 10.10 from Piketty 2014). For more on sources and series, see <http://piketty.pse.ens.fr/capital21c>.

Note: The rate of return to capital (after tax and capital losses) fell below the growth rate during the 20th century, and may again surpass it in the 21st century.

labor earnings, tastes for saving and bequests, and so on. Conversely, a reduction of the variance of these shocks could lead to a decline in wealth inequality. The gap between $(1 - t)r$ and g is certainly not the only determinant of steady-state wealth inequality. It is one important determinant, however, and there are reasons which might push toward a persistently high gap between the net-of-tax rate of return $(1 - t)r$ and the growth rate g in the 21st century—which might in turn lead to higher steady-state wealth inequality (other things equal). In my book, I particularly emphasize the following three potential forces: global tax competition to attract capital; growth slowdown and technical change; and unequal access to high financial returns (Chapters 10–12). Here, I restate and sharpen some of the main arguments.

As international competition intensifies to attract investment, it is plausible that capital taxes will fall, as they have already been doing in many countries in the last few decades. By capital taxes, I include both corporate profit taxation and wealth and inheritance taxes. But of course, the ultimate effect of tax competition will depend on the institutional response. If a sufficiently large number of countries manage to better coordinate to establish a common corporate tax on large corporations and a reliable system of automatic transmission of information of

cross-border financial assets, then the effective capital tax rate might rise, in which case $(1 - t)r$ will decline, and so will steady-state wealth inequality. Ultimately, the outcome depends on the institutional response. Indeed, recent research indicates that better international fiscal coordination is difficult but by no means impossible (Zucman 2014).

Note also that a decline in capital tax rates and a rise in the after-tax rate of return $(1 - t)r$ might in principle induce an increase in saving rates and capital accumulation, thereby leading to a decline in the marginal product of capital which could partly undo the rise in the after-tax rate of return. Indeed, in the example mentioned earlier of the benchmark infinite-horizon dynamic model with no shock and a representative agent, in the long run, the after-tax rate of return to capital has to follow the rule $(1 - t)r = \theta + \gamma g$. In this case, the tax cut leads to a savings response that ultimately moves the rate of return completely back to its earlier level. However, this outcome only arises due to an extreme and unrealistic assumption: namely, the long-run elasticity of saving and capital accumulation with respect to after-tax rate of return is infinite in such a model. In more realistic dynamic models of capital accumulation where this elasticity is positive but not infinite, a decline in capital tax will lead to a net increase in the after-tax rate of return in the long run.⁶

The effect of a growth slowdown on $r - g$ and on the long-run dynamics of wealth inequality is more complicated to analyze. In the historical data, the pre-tax rate of return r seems to display little historical variation, so that $r - g$ definitely appears to be smaller than when the growth rate is higher, as illustrated earlier in Figure 1. This would tend to support the view that lower growth rates in the 21st century (in particular due to the projected decline of population growth) are likely to contribute to a rise of $r - g$.⁷

From a theoretical perspective, however, the effect of a decline in the growth rate g on the gap $r - g$ is ambiguous: it could go either way, depending on how a change in g affects the long-run rate of return r . This depends on a mixture of forces, including saving behavior, multisector technological substitution, bargaining power, and institutions. Let me summarize the main arguments (see Chapters 5–6 of my book for a more thorough analysis; see also the discussion of this point by Jones in this symposium). Generally speaking, a lower g , due either to a slowdown of population and/or productivity growth, tends to lead to a higher steady-state capital–output ratio $\beta = K/Y$, and therefore to lower rates of return to

⁶ For a class of dynamic capital accumulation models with finite long-run elasticities of saving with respect to after-tax rates of returns, and for a study of corresponding socially optimal tax rates on capital, see Piketty and Saez (2013). One of the important findings is that the optimal tax rate is an increasing function of $r - g$ (due in particular to the inequality effect of $r - g$).

⁷ This conclusion largely depends on the way the corrected rates of return reported on Figure 1 were constructed: specifically, the rates of return implied by conventionally measured capital shares are generally very large in high-growth, reconstruction periods. Chapter 6 of my book offers a discussion as to why such high returns might include substantial entrepreneurial labor input and should therefore be corrected downwards; such corrections are highly uncertain, however.

capital r (for given technology). The key question is whether the fall in r is smaller or larger than the fall in g . There are, in my view, good reasons to believe that r might fall less than the fall in g , but this issue is a complex one.

In the benchmark dynamic model, the steady-state β rises as g declines, and the rate of return $r = \theta + \gamma g$ drops. Whether $r - g$ rises or declines as g declines depends entirely on whether the curvature of the utility function γ is smaller or larger than one. However this model does not seem to be particularly realistic empirically, so this may not be the best way to look at the problem. Note that the dynamic model can be viewed as a special case of the general Harrod–Domar–Solow steady-state formula $\beta = s/g$. In effect, in the steady-state of the dynamic model, the (net-of-depreciation) saving rate $s = s(g)$ rises moderately with g , so that $\beta = s(g)/g$ is a declining function of g .⁸

If one instead assumes a fixed, exogenous saving rate s , then the steady-state capital output ratio $\beta = s/g$ will rise even more strongly as g declines. With perfect competition and a constant-elasticity-of-scale production function, whether the resulting decline in r will more than compensate for a decline in g depends (among other things) on the value of the elasticity of substitution. With high substitutability between capital and labor (which might happen because of the rise of new capital-intensive technologies such as robots of various sorts), the rate of return will decline relatively little as β rises, so that $r - g$ will be higher with lower g .⁹ In recent decades, the rise in the capital–income ratio β came together with a rise in the net-of-depreciation capital share α , which in a one-good model with perfect competition implies an elasticity of substitution higher than one. However, the one-good, perfect competition model is not a very satisfactory model, to say the least. In practice, the right model to think about rising capital–income ratios and capital shares is a multisector model (with a large role played by capital-intensive sectors such as real estate and energy, and substantial movements in *relative* prices) with important variations in bargaining power over time (see Chapters 5–6; see also Karababounis and Neiman 2014 about the role played by the declining relative price of equipment). In particular, intersectoral elasticities of substitution combining supply and demand forces can arguably be much higher than within-sector capital–labor elasticities.

Note also there is, of course, no reason why the net-of-depreciation saving rates should be viewed as a constant. What I have in mind is an intermediate model (intermediate between the dynamic model and the exogenous saving model), with a relatively low elasticity of saving behavior with respect to r over a large range of middle returns (say, from 3 to 6 percent) and a much higher elasticity if rates of return take very low or very high values. In particular, if g becomes increasingly

⁸ With a Cobb–Douglas production function $Y = F(K, L) = K^\alpha L^{1-\alpha}$, the long-run capital–output ratio is given by $\beta = \alpha/r = \alpha/(\theta + \gamma g) = s(g)/g$, with $s(g) = \alpha g/r = \alpha g/(\theta + \gamma g)$. See Piketty and Zucman (2014).

⁹ With $Y = F(K, L) = [aK^{(\sigma-1)/\sigma} + (1 - a)L^{(\sigma-1)/\sigma}]^{\sigma/(\sigma-1)}$, the marginal productivity of capital is given by: $r = F_K = a(Y/K)^{1/\sigma} = a\beta^{-1/\sigma}$.

close to zero, then it is clear that $\beta = s/g$ will not go to infinity: otherwise the rate of return would go to zero, and most agents would probably stop saving. In historical periods with very low growth rates (such as in pre-industrial societies), we observe large capital-income ratios, but not infinite β . As pointed out by Jones (in this symposium) and others, another obvious reason why β will not go to infinity is that depreciation would then become enormous. This intermediate model might explain why the rate of return seems to display limited systematic variations in the long run: it is roughly stable within a given range, which one might interpret as an interval of psychologically plausible time preference parameters.

Yet another way to explain why the rate of return appears to be relatively stable in the long run is the following. Pure economic reasoning tends to imply that higher growth leads to higher returns. But high growth periods arguably require more entrepreneurial labor in order to reallocate capital continually and thus to benefit from higher returns (in other words, measured rates of return must be corrected downwards in order to take into account mismeasured labor input in high-growth societies). Conversely, measured rates of returns might be closer to pure returns in low-growth societies (where it is relatively easier to be a rentier, since capital reallocation requires less attention). This is the interpretation that I favor in the book; indeed, the historical estimates of rates of return in the book (those given above and in Chapter 6 of the book) are largely built upon this assumption.

If we combine all these different effects, it is clear however that there is no general, universal reason why $r - g$ should increase as g declines: it could potentially go either way. Historical evidence and new technological developments suggest that it should increase (and I tend to favor this conclusion), but I fully agree that this remains relatively uncertain.

Finally, the last reason (and arguably the most important one) why $r - g$ might be high in the 21st century is due to unequal access to high financial returns. That is, even though the gap between the average rate of return r and the growth rate g is not particularly high, it could be that large potential financial portfolios have access to substantially higher returns than smaller ones. In the book, I present evidence suggesting that financial deregulation might have contributed to such an evolution (Chapter 12). For example, according to *Forbes* rankings, the wealth of top global billionaires seem to be rising much faster than average wealth, as shown in Table 1. This evolution cannot continue for too long, unless one is ready to accept an enormous increase in the share of world wealth belonging to billionaires (and a corresponding decline in the share going to the middle class). Also, larger university endowments tend to obtain substantially higher returns, as shown in Table 2 (and the data presented by Saez and Zucman 2014 on nonprofit foundations indicates a similar pattern). This data is clearly imperfect and too incomplete to prove the general theme of unequal access to high returns. But given that even small changes in $r - g$ can have large amplifying effects on changes in wealth inequality, this effect is potentially important.

Overall, there remains substantial uncertainty about how far wealth inequality might rise in the 21st century, and we need more transparency and better

Table 1
The Growth Rate of Top Global Wealth, 1987–2013

	<i>Average real growth rate per year (after deduction of inflation) 1987–2013</i>
For top 1/(100 million) highest wealth-holders (about 30 adults out of 3 billion in 1980s, and 45 adults out of 4.5 billion in 2010s)	6.8%
For top 1/(20 million) highest wealth-holders (about 150 adults out of 3 billion in 1980s, and 225 adults out of 4.5 billion in 2010s)	6.4%
For average world wealth per adult	2.1%
For average world income per adult	1.4%
For world adult population	1.9%
For world GDP	3.3%

Source: Table 12.1 from Piketty (2014). For more information, see <http://piketty.pse.ens.fr/capital21c>.

Notes: Between 1987 and 2013, the highest global wealth fractiles have grown at 6–7 percent per year, versus 2.1 percent for average world wealth and 1.4 percent for average world income. All growth rates are net of inflation (2.3 percent per year between 1987 and 2013).

Table 2
The Return on the Capital Endowments of US Universities, 1980–2010

	<i>Average real annual rate of return (after deduction of inflation and all administrative costs and financial fees) 1980–2010</i>
For all universities (850)	8.2%
Harvard-Yale-Princeton	10.2%
Endowments higher than 1 billion \$ (60)	8.8%
Endowments between 500 million and 1 billion \$ (66)	7.8%
Endowments between 100 and 500 million \$ (226)	7.1%
Endowments less than 100 million \$ (498)	6.2%

Source: Table 12.2 from Piketty (2014). For more information, see <http://piketty.pse.ens.fr/capital21c>.

Notes: Between 1980 and 2010, US universities earned an average real return of 8.2 percent on their capital endowments, and even more for higher endowments. All returns reported here are net of inflation (2.4 percent per year between 1980 and 2010) and of all administrative costs and financial fees.

information about wealth dynamics. In my view, one main benefit of a progressive wealth tax is that it would produce better information regarding the size and evolution of different wealth groups such that the wealth tax could be adapted in the future on the basis of this better information. I agree with the argument by Kopczuk in this symposium that the data sources about the distribution of wealth that we have at our disposal are insufficient. At this stage, however, it seems to me that the method that infers wealth from the resulting income flows, the income capitalization method developed by Saez and Zucman (2014), produces probably the most reliable estimates we have, and these estimates show substantial recent rise in US wealth inequality—indeed, a higher rise than what I report in my book. In particular, Saez and Zucman find increasing concentration of capital income for all asset income categories (including dividend and interest, which cannot easily be contaminated by labor income considerations). Finally, the Saez and Zucman findings are consistent with the finding from the *Forbes* rankings that the wealth of top wealth-holders is rising much faster than average wealth. However, it is clear that these evolutions remain relatively uncertain. In my view, this makes the lack of transparency about wealth dynamics—largely due to the absence of a comprehensive wealth tax and the limitations of international coordination—particularly problematic.

Toward a New Historical and Political Economy Approach to Institutions

In my book *Capital in the Twenty-First Century*, I attempt to develop a new historical and political economy approach to the study of institutions and inequality dynamics. Economic forces such as the supply and demand for skills, wage bargaining models, or the effect of $r - g$ on wealth dynamics, also play a role. But ultimately, what really matters is the interaction between economic forces and institutional responses, particularly in the area of educational, labor, and fiscal institutions. Given my strong emphasis on how institutions and public policies shape the dynamics of income and wealth inequality, it is somewhat surprising that Acemoglu and Robinson argue in their contribution to this symposium that I neglect the role of institutions. It seems to me that we disagree less intensively than what they appear to believe, and that the well-known academic tendency to maximize product differentiation might be at work here.

It is also possible that some of the confusion comes from the fact that we do not have exactly the same approach to the study of “institutions.” However I believe that our approaches are broadly consistent and complementary to one another: they differ in terms of specific institutional content, as well as in time and geographical scope, more than in substance. In some of their earlier work, Acemoglu and Robinson mostly focused upon a relatively specific institution, namely the protection of property rights. In their fascinating book *Why Nations Fail*, they develop a broader view of institutions and stress the distinction between “inclusive” and “extractive”

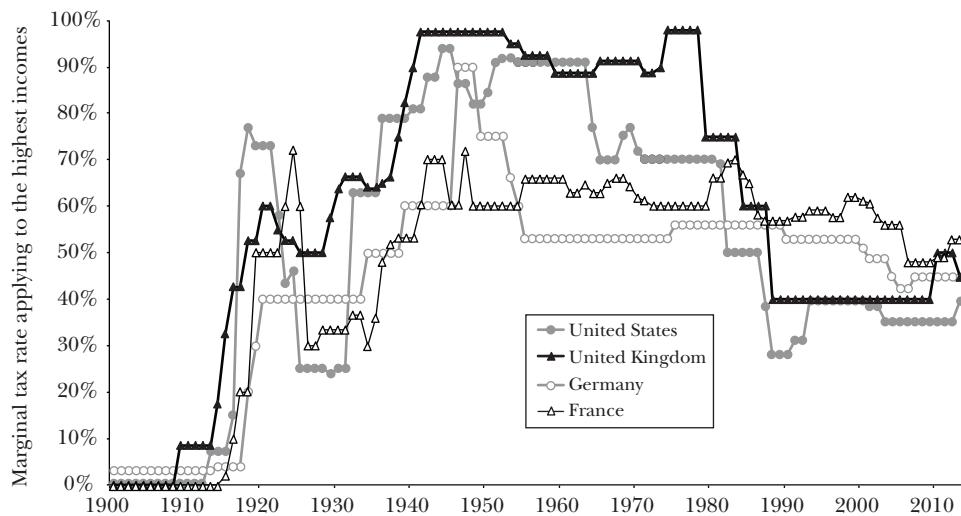
institutions. This broad concept might certainly include the type of institutions and policies on which I focus upon, including progressive taxation of income, wealth, and inheritance, or the modern welfare state. I must confess, however, that seeking to categorize institutions with broad terms like these strikes me as maybe a little too abstract, imprecise, and ahistorical.

I believe that institutions like the welfare state, free education, or progressive taxation, or the effects of World War I, the Bolshevik revolution, or World War II on inequality dynamics and institutional change, each need to be analyzed in a precise and concrete manner within the historical, social, and political context in which they develop. While Acemoglu and Robinson (2012) in their earlier book take a very long-run perspective on the history of the planet (from prehistoric times to the “great discoveries” and the formation of the modern world), I tend to focus on the historical periods and countries on which I was able to collect systematic data, that is, on the 18th, 19th, and especially the 20th centuries (an important period indeed for the formation of the modern social and fiscal state).

My approach to institutions emphasizes the role of political conflict in relation to inequality. In particular, wars and revolutions play a large role in my account of inequality dynamics and institutional change in the 20th century. Of course, steady democratic forces caused by the extension of suffrage also played an important role in the rise of more inclusive social, educational, and fiscal institutions during the 19th and 20th centuries. But many of the most important changes did not come simply from the steady forces of peaceful electoral democracy: rather, specific historical events and political shocks often played an important role. For example, there is little evidence of a natural movement toward more progressive taxation until the violent military, political, and ideological shocks induced by World War I (see Figure 3). Belief systems and collective representations about social inequality and the role of government were deeply affected by World War I and the rise of communism, as they were by the Great Depression, World War II, and then, at the end of the 20th century, by the stagflation of the 1970s and the fall of the Soviet Union.

It is particularly interesting to note that until 1914, the French elite often justified its strong opposition to the creation of a progressive income tax by referring to the principles of the French Revolution. In the view of these elites, France had become equal after 1789 thanks to the end of aristocratic privileges and the development of well-protected property rights for the entire population. Because everybody had been made equal in their ability to hold property, there was no need for progressive taxation (which would be suitable for aristocratic Britain, the story went, but not for republican France). What I find particularly striking in this pre-1914 debate is the combination of strong beliefs in property-rights-centered institutions and an equally strong denial of high inequality. In my book, I try to understand what we can learn from the fact that wealth inequality was as large in France in 1914 as in 1789, and also from the fact that much of the elite was trying to deny this. I believe there are important implications for the current rise in wealth and income inequality and the current attempts to minimize or deny that they are occurring. Then as now, when various shocks are tending to push wealth (and income) inequality higher at

Figure 3
Top Income Tax Rates, 1900–2013



Source: Author (figure 14.1 from Piketty 2014). For more on sources and series, see <http://piketty.pse.ens.fr/capital21c>.

Note: The top marginal tax rate of the income tax (applying to the highest incomes) in the United States dropped from 70 percent in 1980 to 28 percent in 1988.

a time when $r - g$ is at sustained high levels, the result can be a concentration of wealth that is high in historical terms.

Of course, I am not arguing that it will always take wars, revolutions, and other disruptive or violent political shocks to make institutional changes happen. In the case of early 20th century Europe, one can certainly argue that extreme inequality contributed to the high social tensions of the time and the rise of nationalism. But beliefs systems and resulting perceptions and policies can also be affected by peaceful public discussion. However we should not take this for granted. It is important to recognize the role of political conflict in the history of inequality and institutional change. It often took major fights to deliver change in the past, and it is not impossible that it will be the same in the future.

More generally, one of the lessons that I draw from this work is that the study of inequality dynamics and institutional change are intimately related. The development of stable institutions and the construction of a legitimate and centralized government are closely linked to the way different societies are able to address the issue of social inequality in a peaceful and orderly manner. In order to put institutions back at the center of economics, I believe that it is also necessary to put the study of distribution back at the center of economics. Institutions do not arise out of harmonious societies populated by representative agents; they arise out of unequal societies and out of conflict. This is again an issue on which the approaches developed by Acemoglu and Robinson and myself are broadly consistent and complementary.

Finally, let me conclude by making clear that my historical and political approach to inequality and institutions should be viewed as highly exploratory and incomplete. In particular, I suspect that new social movements and political mobilizations will give rise to institutional change in the future, but I do not pursue this analysis much further. As I look back at my discussion of future policy proposals in the book, I may have devoted too much attention to progressive capital taxation and too little attention to a number of institutional evolutions that could prove equally important, such as the development of alternative forms of property arrangements and participatory governance. One central reason why progressive capital taxation is important is that it can also bring increased transparency about company assets and accounts. In turn, increased financial transparency can help to develop new forms of governance; for instance, it can facilitate more worker involvement in company boards. But these other institutions also need to be analyzed on their own terms.

The last chapter of my book concludes: “Without real accounting and financial transparency and sharing of information, there can be no economic democracy. Conversely, without a real right to intervene in corporate decision-making (including seats for workers on the company’s board of directors), transparency is of little use. Information must support democratic institutions; it is not an end in itself. If democracy is someday to regain control of capitalism, it must start by recognizing that the concrete institutions in which democracy and capitalism are embodied need to be reinvented again and again” (p. 570). I do not push this line of investigation much further, which is certainly one of the major shortcomings of my work. Together with the fact that we still have too little data on historical and current patterns of income and wealth, these are key reasons why my book is at best an introduction to the study of capital in the 21st century.

References

Acemoglu, Daron, and James A. Robinson. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Publishers.

Atkinson, Anthony B., and Alan J. Harrison. 1978. *Distribution of Personal Wealth in Britain*. Cambridge University Press.

Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. 2011. “Top Incomes in the Long Run of History.” *Journal of Economic Literature* 49(1): 3–71.

Champernowne, D. G. 1953. “A Model of Income Distribution.” *Economic Journal* 63(250): 318–51.

Karabarbounis, Loukas, and Brent Neiman. 2014. “Capital Depreciation and Labor Shares around the World: Measurement and Implications.” NBER Working Paper 20606.

Kuznets, Simon. 1953. *Shares of Upper Income Groups in Income and Savings*. National Bureau of Economic Research.

Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Harvard University Press.

Piketty, Thomas, Gilles Postel-Vinay, and Jean-Laurent Rosenthal. 2006. “Wealth Concentration in a Developing Economy: Paris and France, 1807–1994.” *American Economic Review* 96(1): 236–56.

Piketty, Thomas, Gilles Postel-Vinay and Jean-Laurent Rosenthal. 2014. "Inherited vs. Self-Made Wealth: Theory and Evidence from a Rentier Society (1872–1927)." *Explorations in Economic History* 51(1): 21–40.

Piketty, Thomas, and Emmanuel Saez. 2013. "A Theory of Optimal Inheritance Taxation." *Econometrica* 81(5): 1851–86.

Piketty, Thomas, and Emmanuel Saez. 2014. "Inequality in the Long Run." *Science* 344(6186): 838–43.

Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva. 2014. "Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities." *American Economic Journal: Economic Policy* 6(1): 230–71.

Piketty, Thomas, and Gabriel Zucman. 2014. "Capital is Back: Wealth–Income Ratios in Rich Countries, 1700–2010." *Quarterly Journal of Economics* 129(3): 1255–1310.

Piketty, Thomas, and Gabriel Zucman. 2015. "Wealth and Inheritance in the Long Run." In *Handbook of Income Distribution*, vol. 2, edited by A. Atkinson and F. Bourguignon, pp. 1303–68. Elsevier.

Saez, Emmanuel, and Gabriel Zucman. 2014. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." NBER Working Paper 20625.

Stiglitz, Joseph E. 1969. "Distribution of Income and Wealth Among Individuals." *Econometrica* 37(3): 382–97.

Zucman, Gabriel. 2014. "Taxing Across Borders: Tracking Personal Wealth and Corporate Profits." *Journal of Economic Perspectives* 28(4): 121–48.

Copyright of Journal of Economic Perspectives is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.