

1. How do we measure unequal? The who, where, what, when, and how of inequality

Before we get into the big “why” issues, we first have to agree about exactly what we are talking about when we talk about inequality. How can we quantify inequality? Talking about data is never the sexiest topic in economics, but the fact of the matter is that obtaining reliable, accurate, and well-defined data is one of the biggest challenges in understanding economic inequality. There are many different types of inequality that we could talk about, and many different sources of data that could be used to measure inequality. A great deal of the public debate surrounding inequality is often at cross-purposes because different people have different ideas about what constitutes inequality and how they choose to measure it. Many people believe in equality as a *principle*, but there is often little agreement about *what* should be made equal.

You can think of this opening chapter as a deep breath before we dive in and immerse ourselves in the big and challenging questions related to the causes and consequences of economic inequality. In this chapter, we examine the many ways that economic inequality can be defined and how we can bring data—or sometimes fail to bring data—to each of these definitions.

In general, confusion about how we measure inequality generally centers around five eternal questions that shape any empirical debate in the social sciences: Who are we talking about? Where are we talking about? What are talking about? When are we talking about? And how are we going to talk about it?

WHO ARE WE TALKING ABOUT?

The most basic “who” question relates to whether we are trying to measure inequality across individuals or households. There is often no clear answer as to whether equality should be measured across individuals or households if we want a clear answer to the question “how unequal are we?”

There are different ways to measure inequality across individuals. Measuring inequality across individual workers tells us a lot about wage and salary inequality that relates directly to our jobs. It can also tell us how much inequality

can be attributed to worker productivity, to differences in earnings, and to individual professions. On the other hand, measuring inequality across all individuals, regardless of whether they work or not, gives us a clearer picture of welfare, or the resources available to each person whether they work or not.

However, most measures of inequality take place across households. Measuring inequality across households tells us more about welfare and standards of living because people's standards of living are not just a function of how much money they earn, but also who else is earning money in the family and how many people need to share each paycheck. Of course, measuring household income creates a number of complications. First, households differ in size. In my house, when my family and I go out to eat I have to multiply menu prices by 6. For other households, that might be 2 or 4 or 8. As a result, economists usually standardize household income to account for the number of family members (called *equivalency scaling*). However, it is important to remember that differences in household size are not simply due to differences in the number of children. There has been a significant increase in single-parent households; in the US, only 6 percent of children were raised in single-parent households in the 1960s, but by 2016 this had risen to 27 percent of children.¹ This has had a dramatic impact on household inequality for obvious reasons—such as fewer households with two parents working—and less obvious reasons related to child-rearing and education that we will discuss later. As a result of these differences in household composition, two households might share the same equivalency-scaled income levels but have very different standards of living.

The second complication is that whatever income a household has, the power to spend this income is not distributed equally. In many households, men dominate spending decisions. As a result, it would be incorrect to assume that every male-led household will have the same standards of living as a female-led household. For example, there is empirical evidence that when women have greater power over resources, more money gets spent on children and less on luxuries.²

A third complication is that changes in labor force participation can make it hard to compare household income across time. In the US, only 22 percent of women worked in 1947; today that number is 59 percent. More women working leads to higher household income, but it also leads to more inequality between those households in which women work and those in which they don't. We will also see in a later discussion that the increased number of women entering the labor force has also magnified inequality between households because working, educated women are more likely today to marry working, educated men.

One final complication is that households leave out a lot of people who are not part of a household: people in the military, nursing homes, and those who

are homeless or who are refugees. As a result, these often-forgotten people are also forgotten in our household income data.

In spite of all of these complications, please keep in mind this general rule: When I talk about inequality in income or wealth in this book, unless otherwise stated, I am talking about inequality among households, not individuals. This is because, as we will see shortly, our most reliable data on incomes comes from sources in which only households report, not individuals.

A second aspect of the “who” question relates to categorical differences related to race and gender. Race is determined by ancestry: it is primarily socially and culturally determined, not genetically determined. While it is only a social construct, race remains one of the most important and enduring factors in creating economic inequality. To pick just a couple of data points from the US, in 1963, the average white family earned \$43,000 more than the average African American family (in constant dollars); by 2013 this difference has risen to \$123,000.³ In 2009, the typical black household holds wealth of \$5,667—barely one-twentieth of the wealth of a typical white household. Hispanic households only hold about one-eighteenth of the wealth of a typical white household.⁴ Home ownership, employment rates, and life-expectancy are lower at every level of education for blacks and Hispanics than for whites.

Gender inequality is also very real. Women tend to work more hours (formal and informal) than men and are paid less for it. In the US, women earn only about 78 percent of what men earn. While some of this difference can be explained by things such as the risk of the job (men tend to work in more risky jobs), a number of jobs have been “feminized” through social norms that lead to implicit or explicit wage discrimination against women. However, the wage gap between men and women has shrunk in most countries since 1960, and almost all of this improvement can be explained by an increase in female educational attainment; today women outnumber men in college within 29 of 32 developed countries.⁵ However, gender inequality exists, even at the top of the income distribution. Only 13 percent of the highest earners in the US are women, and only 3 percent of the world’s billionaire wealth is held by self-made billionaire women.⁶ In fact, in the US there are more male CEOs named John than there are women CEOs (and the same holds for David as well).⁷

One final “who” question that must be answered relates to membership: Who are we including and who are we excluding? When we talk about inequality on a national level, we need to be clear who gets counted. Are we including only citizens? Are we including recognized refugees or those with legal residence without citizenship? Or are we including all residents? These are difficult questions to answer because not everyone who has legal residence or even citizenship in one country has the same rights as other citizens in other countries. For example, the *Houku* system in China severely restricts the

ability of those born in rural areas to migrate and enjoy public services in many urban areas. Likewise, many external immigrants can live in many developed countries but with severely restricted economic and civil rights. In these cases where people are purposefully treated differently, it is not clear that we should treat everybody the same in our inequality statistics. In other words, a country with a great deal of *income inequality* but that has severe *rights inequality* may not necessarily be more equal in terms of quality of life. However, these differences are simply not captured in most of our economic inequality data, as most of our data is collected from all citizens without regard for others who live in the country but under a different legal status.

When we talk about differences across race, gender, or citizenship status, we cannot avoid the topic of discrimination—that people are purposefully treated differently because of their categorical status, and that belonging to a category conveys privileges and costs, either consciously or implicitly. As we will discuss later in the book, there is a strong presumption in economics that discrimination is inefficient, and as a result that it should not persist in a market economy because firms that discriminate will be forced out of competitive markets. However, it is hard to reconcile efficiency and perfectly competitive markets with a growing body of research that finds that persistent discrimination is a fact of life. In fact, we will see that discrimination may be inefficient but still persists because of psychological biases and the fact that markets are not perfectly competitive. To highlight just one study of discrimination in the job market, white male job applicants with a criminal record are more likely to be called back for a job interview than black men without a criminal record.⁸

So be warned: In the interest of conciseness, our inequality data will often aggregate across people and households and talk about inequality in the US, or inequality in Britain, or even global inequality. But these summary statistics about inequality will themselves hide a great deal of inequality across groups of individuals. I will keep this in mind as this book proceeds, and I hope that you will as well.

WHERE ARE WE TALKING ABOUT?

One of the most underreported facts in inequality is that geography plays the primary role in who has what. For reasons that I will elucidate throughout this book, the three most important factors in inequality are the same as the three most important factors in the real estate business: location, location, and location. In talking about inequality, *where* we are talking about is more important than *who* we are talking about.

Where matters in many different ways. When I refer to *domestic*, or *within-country inequality*, I am talking about national inequality within the US, within the UK, or within China, for example. When talking about

within-country inequality, it is domestic households led by citizens (not just simply residents in most data) that are the category of interest. On a day-to-day basis, when we see that our neighbor's house is bigger than ours or we hear that Kanye West and Kim Kardashian bought a gold toilet seat (I wish I made this up), it is within-country inequality that is foremost in our minds, because this is where we live our lives.⁹

However, levels of domestic inequality within countries, as I will explain in the next chapter, are dwarfed by *between-country inequality*, or the inequality across different countries. For example, per capita income in the US is more than 180 times larger than it is in the Republic of the Congo.

Finally, we will also examine *global inequality*, which ignores nationality and compares households (to the extent that the data allows) across the globe. As a result, global inequality essentially combines within-country inequality and between-country inequality.

Thinking hard about the differences in within, between, and global inequality gives us a more nuanced view of what is happening to inequality across the world. We will see that between-country inequality appears to be getting worse if we just look at the incomes of the richest countries compared to the incomes of the poorest countries. But if we take into the account the population size and rapid growth of China and India, then between-country inequality is actually declining. As a result, even with within-country inequality rising in most countries today global inequality is still falling. This seeming contradiction—that inequality is rising inside of most countries but declining across the world's population—lies at the heart of what makes this such a confusing economic era.

One more important statistical issue related to location must be addressed: Whenever we talk about where, we have to not only consider the fact that incomes are different across the globe, but currencies and prices are also different. We not only need to adjust for exchange rates—which is relatively easy given that there is market data for this—but we also need to adjust for the fact that poorer countries generally have lower prices and, as a result, households get more purchasing power for a given level of income. For example, five US dollars will buy you more food in Nigeria than in Japan. Thankfully, economists have put together the International Price Comparison Project (ICP), which is the largest empirical study in the history of economics.¹⁰ The ICP collects price data at regular intervals in all countries in order to construct what are called *purchasing power parity (PPP) exchange rates*. PPP exchange rates are the exchange rates at which, in theory, a person could buy the same amount of goods across countries with the same amount of income. Market exchange rates can be very different than PPP exchange rates because the prices of similar goods and services can be very different across countries. To pick one example, in 2011 the market exchange rate averaged 46 Indian rupees

for every US dollar. However, the PPP exchange rate as measured using ICP data was only 15 rupees per US dollar. In other words, you get three times the purchasing power in reality than the market exchange rate for rupees would suggest; or, to say this in another way, Indians actually have three times the standard of living they would appear to have if you used unadjusted market exchange rates. This is a big deal, which is why using PPP exchange rates is absolutely essential to having any sort of accurate measure of between-country or global inequality.¹¹

Of course, there are problems with the PPP exchange rate data that we use. First, the ICP compares the prices of similar bundles of goods across countries. On one hand, assuming that consumption baskets are the same across countries makes it easy to compare prices. But on the other hand, the goods that are typically consumed in countries can be very different. If Americans, Japanese, and South Africans buy very different things, how are we to compare their relative costs of living? Likewise, even people within the same country don't pay the same prices and purchase the same things. In some countries, the differences between prices in urban and rural areas are at least as great as they are across borders; Shanghai's prices are closer to those of New York than they are to a rural village in the far west of China. The ICP makes some technical adjustments for these problems, but they are fraught with ambiguity. The point here is not to say that PPP exchange rates are biased or uninformative, but instead to emphasize that they must be used with cautious skepticism. Every piece of data related to economic inequality is imperfect; however, this is not the same thing as saying that the data is uninformative.

The urban versus rural divide is not just important in regards to differences in prices, it is also real in regards to differences in income levels. In China, urban incomes are three times those of rural incomes in 2015, up from only twice as large in 1983; these regional differences are the largest factor in the overall rise in inequality in China.¹² In the US, a significant portion of the rise in income inequality in the late 1990s was driven by rising incomes in only five counties: Manhattan (New York), King County (Seattle), and the three counties that comprise Silicon Valley (California).¹³ Likewise, regional inequality has also been driven by high concentrations of firms in specific industries that are either doing well, or not so well. In the US, when the finance and IT sectors boomed, Manhattan and Silicon Valley boomed as well, while the downturn in the automobile and heavy manufacturing industries led to busts in Detroit and other Midwestern rustbelt cities. It is rising regional inequality in developed countries like the US and the UK that has driven the big divide in region voting patterns seen in these two countries during the 2016 US presidential election and the UK Brexit vote.

WHAT ARE WE TALKING ABOUT?

When we talk about what is unequal, the most commonly used data used to measure economic inequality are *pay (earnings) data*, *income data*, *wealth data*, and *consumption data*. Let's talk about each of these and how they differ.

Pay, or *earnings*, refers to compensation for work. Pay data incorporates wages and salaries, but excludes capital income from savings and corporate profits. It is data that is associated with a job, not necessarily a particular person. So we can talk about the earnings of elementary teachers, software engineers, and CEOs. Pay data is the easiest data to collect in rich countries because employers (particularly large ones) are easier to track over time, there are fewer firms than households, and they often have administrative obligations for reporting such data. The value of pay data is in its reliability and comparability across jobs, despite the fact that there can be differences in how firms define hours worked or enforce labor regulations. Most importantly, pay data allows us to measure one important cause of increasing economic inequality: Earnings have grown more unequal between different jobs, and even among people doing the same job at different firms. This is something that we will examine as we talk about the causes of increasing inequality today.

Unfortunately, pay data also has two significant limitations. First, it ignores capital income from savings and investments, or other sources of income that households may be able to collect, making it much too narrow a measure of overall economic inequality. Capital gains, dividends, stock options, interest, and rental income are all ignored in pay. A second problem is that pay data is quite limited in many poorer countries where many workers work on an informal basis (40 to 60 percent of all workers in many less-developed countries), meaning that most employees are missed in any pay data collected only from registered firms.

Income is a flow variable, meaning it is measured over a period of time, typically a year. Income is broader than pay because it includes income from work and income from capital (stocks, bonds, trusts, rents). To give you an idea about how important capital income is in determining overall income inequality, consider data from the US on capital gains, which is income received from the appreciation of an asset such as the sale of a stock, bond, or real estate. If capital gains were excluded from income, the top 1 percent would go from receiving 22.5 percent of all income to 19.3 percent of all income.¹⁴

The best way to collect income data would be from income tax records. Income tax data is more reliable—because there are potential penalties for falsifying it—and it gives us the best data on those with the highest incomes. It would be nearly impossible to get accurate income estimates of the richest people in the world without income tax data. The problem is that while income

tax data is available for the US, many other countries, particularly poorer countries, don't have income taxes (they pay consumption taxes) or income taxes only apply to a small portion of the population. Complete income tax data only exists for 29 countries, and comparing data across these countries and even within the same country over time is very difficult because of different and changing tax laws. There is also the problem that people, but particularly the rich, hide income. They do this in legal ways, through deductions, loopholes, exclusions, offshoring, or by changing their behavior to allow them to avoid taxes. They also do this in illegal ways as well, through tax evasions such as underreporting or manipulating stated income.

Survey data of households is popular when good income tax data is not available—it is probably our best source of income data in most countries. The most common of these are International Household Income surveys. These surveys are conducted across countries and ask the heads of households to specify the range that their income falls into—for example, \$10,000–\$20,000, \$20,000–\$30,000, etc., all the way up to a truncated top income of \$999,999. However, this survey income data still has flaws. The definitions of what counts as income is not standardized across countries. Also, it is clear that people may not fill these surveys out accurately or truthfully—this is likely to be a particularly big problem among the richest households, who have more complicated financial pictures and are often reluctant to report income they might have hidden away (or offshored) from tax authorities. However, surveys also miss poorer households (who might be homeless, for example) and miss important sources of income in poorer households, particularly “in-kind”, or barter income. As a result, these surveys underreport the very rich and the very poor, meaning that they understate actual income inequality by as much as 10 percent.¹⁵ The final problem with income survey data is that sample sizes are limited, and participation is biased. In 2013, 200,000 surveys were distributed in the US, 100,000 in India, but only 10,000 in most countries. Many surveys are not returned, and participation rates have been falling across the globe.¹⁶ In the US, the nonresponse rate was 15 percent in 2017. Based on areas where nonresponse rates are the highest—which was in the richest and poorest neighborhoods—nonparticipation in the survey may lead to 10 percent underestimation of actual income inequality.¹⁷

Because of rising interest in the topic of inequality, there is more survey data being collected today than ever before. For example, the Luxembourg Income Study is a standardized survey with lots of microeconomic data on individual households, but is only collected in 50 rich countries. The United Nations, World Bank, and many regional organizations and national central banks have been collecting their own income survey data as well. Much of this data is becoming more standardized over time, but the fact is that perfect comparability across countries will never be possible because of differences

in how different countries tax income, pay for healthcare and education, or enforce accounting standards.

There is one more important factor in measuring income: How do we deal with government? The government taxes most households, transfers income to some people (in particular the poor and elderly), and provides public goods such as education and healthcare to everyone. How do we account for all of these government programs when measuring income?

In fact, there are three different kinds of income that we can use to measure inequality. *Market income* is income from market activities: wages, salaries, and capital income. There is no adjustment for government taxes and transfers and, as a result, market income tends to be the most unequal across households.

Gross income is market income plus government transfers, such as public pensions, medical and educational payments, and unemployment insurance. Gross income really looks at the total resources available to a household, but is also somewhat unrepresentative of actual welfare because it ignores the fact that people have to pay for these government benefits.

Finally, there is *disposable income*, which subtracts income tax collections from gross income. Disposable income does not count consumption or sales taxes that people have to pay, and it is important to keep this in mind given that poorer households tend to spend a bigger share of their income on consumption. Disposable income is the best measure of what a household actually has to spend, and generally is the best measure of the actual quality of life of individual households. Disposable income is generally less variable than market or gross income because the tax systems in most countries are more or less progressive, meaning those with higher income pay higher tax rates. However, in poorer countries where government is less effective and more income is earned informally, there are often smaller differences between market, gross, and disposable income. Also, in countries that rely more heavily on consumption taxes, gross income will be closer to disposable income.

Wealth is a stock variable, meaning it is measured at one point of time. Instead of looking at income that is measured over a fixed period of time, wealth looks at the accumulated assets from savings, the returns on investment, and any inheritance or transfers received by households. To calculate wealth, you must assess the value of a household's financial assets (stocks, bonds, real estate, cash), their nonfinancial assets (houses, real estate, art, jewelry, cars), and subtract any liabilities (debt or other borrowings) that they have. Unfortunately, traditional measures of wealth focus on tangible assets and ignore many intangible assets such as the value of education, networks, and other social connections. These intangible assets are increasingly important in explaining inequality, as we will talk about later.

The big advantage of wealth data is that we can measure what is going on with the very rich because most survey data on incomes is truncated at

\$1 million, while wealth data is not truncated in the same way. Wealth data primarily comes from three different data sources, and economists have tried to put these sources together as best as they can. First, wealth data can be obtained from countries that have a wealth tax, such as Sweden. However, wealth taxes do not exist in most countries and tend to be paid by only a small number of people when they do exist. Second, wealth can also be imputed from the capital income that people report to tax authorities or on income surveys. These imputations involve complicated estimation methodologies that are fraught with missing data, modeling assumptions, tax avoidance, and changing tax laws among other problems.¹⁸ Finally, probably the best source of wealth data is from personal wealth surveys. The US Federal Reserve and European Central Bank conduct wealth surveys regularly, although they tend to have high nonresponse rates from the richest people, and it is always unclear how truthful people are willing to be about how much they are really worth. Other common surveys of wealth are found in *Forbes* magazine's "500 Richest People in the World" and *The Sunday Times*' "Rich List". While these look at only a small number of rich people, there is actually good reason to think that the participants in these lists are more truthful, or might even overstate wealth because of the status associated with making these lists.

Wealth is dramatically more unequal than income. Some people are extremely wealthy, while roughly one-third of those living in developed countries such as the US have zero or negative wealth, meaning that their debts are greater than their assets. In fact, half of the world's population owns 100 percent of the wealth, meaning that wealth is not anything close to a global norm. Wealth data alone cannot give us the full picture of inequality because so many people's wealth cannot be accurately measured. But without talking about wealth, it is impossible to get the full picture of economic inequality.

Consumption, in a very real sense, is what matters. We work and save to consume, not just to earn an income. Consumption is the primary determinant of our current welfare and standards of living.

Because people tend to smooth their consumption over time, inequality in consumption tends to be lower than inequality in income, which is much more variable. For example, the children of rich parents have the backstop of their parents' wealth which makes them better able to smooth their income over the course of their life even if they are not earning any income—a backstop that the children of poor people do not have. In this sense, inequality in consumption tracks true inequality better than income inequality.

All of this said, there are problems with using consumption data to measure inequality. First, we collect data on consumption expenditure, not actual consumption. As a result, consumption data (like income data) misses a lot of informal consumption (e.g. growing your own food, barter, trade in household services). There are also difficulties in adjusting for different costs of living in

different regions and among different people that makes getting an accurate picture of consumption across households difficult. But there are also more fundamental problems associated with using consumption as a measure of inequality. Income gives you more than just the power to consume. Having income gives you the power to control resources, and people with income are empowered in ways that the poor can never be, regardless of their consumption choices. In the words of Amartya Sen, income gives people “capabilities” and options in life to make their own decisions and choose how they would like to participate in society.¹⁹ Two people might have the same consumption levels, but there are decidedly more options for the person who has a high income and chooses to limit their consumption than there is for the person who has little income and has their consumption dictated to them. As a result, income is a better measure of the true constraints created by inequality.

Another reason why income matters more than consumption is that income also reduces risk, and reducing the risks associated with life is just as valuable as maintaining high current levels of consumption. You will not have to be convinced of this if you have ever missed a meal from lack of money, or have not known how you will pay your next school fee, or have ever seen someone you love refused medical treatment because they cannot pay for it. Someone who lives frugally but is still rich will never really know these stresses and how they impact the quality of life.

The good news is that we don’t really have to choose between income and consumption to understand trends in inequality. Consumption inequality closely tracks income inequality, although the growth in consumption inequality has not been as rapid.²⁰

WHEN ARE WE TALKING ABOUT?

Time has to be an element of consideration when examining inequality for a couple of reasons. First, prices and the cost of living rise over time, so we must use *real data* (i.e. adjusted for inflation and changes in the price level) when comparing data over time. Of course, calculating price levels and the cost of living is difficult to do, for the reasons we talked about when we talked about PPP. But even if our estimates of the cost of living are imperfect, they are better than ignoring the problem. In this book, I will always be using real statistics when comparing data from different points in time.

Time also has to be considered because our inequality data are always associated with a specific moment in time—they are photographs, not films. As a result, if we plotted everyone’s incomes from the poorest to the richest, this graph could look exactly the same year after year, but specific people would have moved to different spots on the curve. This often has to do with people’s ages: Many people start out at low levels of income, see their incomes

grow until they near retirement age, then typically see their incomes fall. In other words, individual income tends to be hump-shaped over a lifetime, and demographics such as the average age of the population can play a role in shaping overall income inequality at any one point in time. For example, the baby boom of the 1950s and today's declining birth rates have led to aging populations in most developed countries. If older people are wealthier than younger people (which tends to happen in countries that are growing), then older populations should also have more wealth inequality.

Age is not the only factor that determines where an individual ranks relative to their peers. One of the things we will examine later in this book is how much "churning" is actually occurring across income rankings. When ranking incomes, are people moving up and down the income ladder a great deal over time and over their lives? Or are the people that are born at the bottom and at the top of the income ladder staying at the bottom and the top? In other words, are there high or low levels of *income mobility*? When there is little income mobility, people's stations in life are largely predetermined and their incomes are not necessarily a function of their own hard work and abilities, but of the income class that they were born into. More churning and greater income mobility might be more consistent with a meritocratic society in which people's income class reflects their contributions to society (although "merit" is always a loaded word, for reasons I will discuss later).

Finally, time plays a big role in measuring inequality because the quality of our data has improved over time. As a general rule, the further back in time you go, the weaker is the evidence. Today we have sophisticated income and wealth surveys. In the 1900s, we had to primarily rely on income tax records. In the 1800s and before, most inequality data that exists comes from social tables collected in some countries, regions, and cities that recorded demographic, income, and wealth records by hand. From these incomplete social tables, contemporary scholars can infer broader measures of inequality. Going even further back in time, the best evidence of inequality can be found in anthropological data patterns, such as land ownership in Roman Egypt, variation in house sizes in Aztec Mexico, and inheritance shares in Babylonia.

HOW ARE WE TALKING ABOUT IT?

Collecting data points on income or wealth is fine, but they are just data points. How do we organize this data and summarize it in ways that allow us to understand what this data is actually telling us about economic inequality?

Let's begin by clarifying what we mean by a *distribution*. The easiest way to visualize a distribution is to create a graph that plots the *amount of a resource* (wealth, income) from the lowest to highest levels on the horizontal axis and the *number that possess this level of the resource* on the vertical axis (either

households or individuals). This graph will likely be hump-shaped, meaning that there will be more people in the middle than are very rich or very poor. It will be truncated on the left, because few people will have zero and nobody can have negative amounts of the resource. On the right end, this distribution is likely to have a long tail that stretches far to the right because there will be a small number of very rich people. We would consider a distribution to be more equal if it looked more like a normal distribution, with a large number of people in the middle and few at the poor and the rich ends—in other words, more hump-shaped and symmetrical. A distribution will be more unequal the more that the distribution flattens out, is pushed to the left, and stretches further to the right. Figure 1.1 presents three hypothetical distributions of income across households. Distribution A would be the most equal—there are more households in the middle, and fewer at the extremes. Distribution C would be most unequal, as there are more poor and rich and fewer middle-class households.




Figure 1.1 Hypothetical income distributions

The distributions in Figure 1.1 are hypothetical. In fact, if we were to make this graph using real data from a country like the US, it requires some manipulation of the data to visualize the entire distribution on one graph because those at the highest end of the distribution earn so much more than those at the bottom and the middle of the distribution that a simple plot will place most households at what looks like zero with a long line stretching out along the horizontal axis.

Measures of inequality, then, are really ways to describe the shape of distributions without resorting to a simple, and less informative, plot of the number of households living at various levels. Our measures of inequality are ways to


quantify disparity within groups that are more concrete than using an “eyeball test”.

So, how exactly are we going to characterize the inequality of different distributions? The first—and most important—issue is to clarify whether we are going to measure absolute inequality or relative inequality. *Absolute inequality* is the absolute differences in levels between individuals. For example, if one household makes \$50,000 and another \$20,000, their absolute level of income inequality is \$30,000. For an entire population, a measure of absolute inequality could be the *range* of the data (the highest value minus the lowest value) or a *decile ratio* (the level of the household that just qualified for the top 10 percent minus the level of the richest household in the lowest 10 percent of the distribution). Absolute measures of inequality are intuitive, but they have a number of problems. First, measures of absolute inequality will always rise if an economy is growing and average incomes are rising. For example, in 2008 (in 2005 dollars), the richest 1 percent of the world earned at least \$71,000 per year, but median income across the globe was only \$1,400, and the poorest 10 percent earned less than \$450.²¹ Thus, a 1 percent increase in income for the top 1 percent would be the same size as a 58 percent increase in the income of the bottom 10 percent. If we focus only on measures of absolute inequality, we will always conclude that inequality is getting worse, even if the incomes of the poor are growing much faster than the rich over time.

Relative inequality is a better way to measure reality because, in economists’ minds, we think that everything is relative. People instinctively weigh gains relative to where they started, not from where they have never been. If a poor and rich person both find \$100 on the street, who should be happier for it? The poor person, obviously, because it changes their financial situation by more—in other words, it will be a larger percentage change in their income. Relative measures of inequality recognize that a \$1 income gain to the poor should be weighted more heavily than a \$1 income gain to the rich if the goal is to accurately measure the impact of inequality. If a country doubles the income of their poorest while the incomes of the richest only slightly improve, we should recognize that inequality has fallen in relative terms and that overall well-being has improved, even if measures of absolute inequality have risen.

There are three different measures of relative inequality that we will use throughout this book. The first is the *Gini coefficient*, named after the Italian statistician Corrado Gini. The Gini coefficient summarizes the entire distribution of income, from the very top income to the very lowest income, in a single number between zero and one that measures the extent to which the distribution of any resource across individuals differs from perfect equality and perfect inequality. A Gini of one means one person owns everything, a Gini of zero means that everyone owns the exact same amount. As a result, a higher level of the Gini coefficient means higher inequality. A Gini coefficient reduces

the entire distribution to a single number, so obviously a lot of information is obscured by a Gini coefficient. But the conciseness and simplicity of interpreting a Gini makes it a useful tool in describing inequality.²² To get an idea about what the Gini coefficient looks like in the US, Figure 1.2 presents the Gini coefficients for market, gross, and disposable income since 1979.



Source: Luxembourg Income Study Database (2019).

Figure 1.2 *Gini coefficients for market, gross, and disposable income in the US*

Here we can see that as measured by the Gini coefficient, income inequality in the US has been rising steadily since 1979. For example, market inequality rise from .43 in 1979 to .53 in 2013. The disposable income Gini rose from .31 to .38, which is also a proportionally substantial increase. We also see that market inequality is greater than gross income inequality, which is greater than disposable income inequality, as we would expect.

There are some important things to keep in mind about Gini coefficients so that we use them appropriately and don't misinterpret, or over-interpret, what they can tell us. The Gini does not depend on the size of the population, or whether we are using absolute or relative data, or what currency units the data is measured in. In other words, Ginis are easy to calculate and easy to compare across countries.

A Gini coefficient will stay the same when percentage changes across the distribution are equal, but this means that the dollar-amount changes in income will be unequal, and the difference between the highest and lowest incomes will increase. In other words, it is a measure of relative inequality, not of absolute inequality. As a result, Gini coefficients do tend to paint a more conservative picture of inequality and minimize the huge differences in levels between

the very rich and the very poor. For example, in the US, the 1 percent richest Americans earn at least 158 times the Americans in the bottom 10 percent. Or, to put it in another way, the top 1 percent can earn in a little more than two days what it takes the bottom 10 percent almost a year to earn. These staggering differences in income are somewhat hidden by Gini coefficients, which essentially report an “average” level of inequality among a population. Likewise, the extent of poverty, or people who have absolute income levels that fall below minimum standards of living, are ignored by Gini coefficients. The bottom 10 percent of the world’s population is living on about \$1.90 day (in 2011 dollars).²³ Thinking about the real difficulties of living on so little brings the reality of global poverty home in a way that a Gini coefficient cannot.

One less obvious problem with using Gini coefficients is that having a Gini coefficient value of one is actually not possible in the real world because everybody but one very lucky person would starve. This raises the question of the maximum feasible Gini—what should we consider a “high” level of the Gini if a value of one is impossible? One approach would be to estimate the minimum level of subsistence needed to feed the population, then subtract this amount from total income and then estimate the distribution of this “surplus” resource. For example, if total income was four times the minimum subsistence level of income to keep everyone alive, then the maximum Gini would be .75—in which case, everyone lives at subsistence except for one person, who owns everything else. If total income was twice the minimum subsistence level, the maximum Gini would be .5. This illustrates an important problem with comparing Ginis across countries: Rich countries will tend to have higher maximum Gini coefficients because they will have more surplus income (i.e. more income above subsistence levels). According to estimates, if the minimum subsistence level of income is \$300 per person (in 1990 dollars), a country with a per capita GDP of \$1000 has a maximum Gini of .76, while for the US the maximum Gini is close to one. However, it is probably more appropriate to consider the “socially acceptable minimum income” when making these calculations instead of just what income is needed to keep people alive. In this case, the maximum Gini for rich countries can be much lower—in the US, as low as .6.²⁴ This puts the current Gini coefficient of disposable income in the US of .38 in a much clearer context: it is two-thirds of maximum possible inequality, and rising.

The second measure of relative inequality that we will use in this book is *income shares*. Income shares capture the percentage of the total income owned by a particular group within the overall distribution. Examples would be the income earned by the top 1 percent, or the income shares broken into groups of 10 percent of the population (deciles), or the income shares broken into groups of 20 percent of the population (quintiles). When using income shares, we are trying to better describe the overall shape of the distribution.

Income shares are more complicated to report than Gini coefficients, but have the potential to convey more detailed information on inequality across the entire distribution.

By expressing data in terms of income shares, we are expressing data in terms of percentiles and not levels, meaning that they are not dependent upon the units of currency we are using, so they can be easily compared across countries.

Of course, income shares aggregate data and obscure some important information; data broken into deciles will convey more information than data broken into quintiles (but also are messier to report and comprehend). However, the biggest thing to keep in mind when looking at income shares is that the same people will not necessarily be in the same percentiles across time. Once again, there is churning within the distribution as people move up and down the relative distribution over their lifetime. We will discuss how much churning there is, and why this is important, when we talk about income mobility. For now, it is important to recognize that income share data imposes the condition that for one person to rise to a higher percentile ranking, another person must fall in the percentile rankings. Or, to say this another way, even if everyone in the lowest income share were to see their absolute incomes improve, they could still fall further behind those in higher percentiles in terms of relative income, and we would observe higher overall measured relative inequality.

There is one final measure of inequality that we will use in the book when talking about global inequality, or inequality across the entire world population, which is known as *Theil's statistic*, named after the University of Chicago statistician Henri Theil. Theil's statistic is very complicated to calculate, but relatively easy to understand. It is useful because it can serve as a measure of overall inequality which can also be separated into a measure of how groups within the overall population differ.²⁵ Theil's statistic is a relative measure of inequality, like the Gini coefficient, but is not limited to being between zero and one. In fact, the absolute level of a Theil's T statistic is meaningless by itself unless it can be tracked over time, in which case it can be used as a measure of how inequality is changing.

Here is where Theil's statistic becomes particularly useful. Let's say you have two groups of people in a population; for example, men and women. When using Theil's statistic to measure inequality, it is true that the inequality of the entire population will be equal to the weighted-average of the differences in income among men and the differences in income among women (within-group inequality), and the difference in average incomes between men and women (between-group inequality). As a result, Theil's statistic can be deconstructed into a measure of within-group inequality and between-group inequality that gives us a much fuller picture of what is going on with the overall distribution.

Theil's statistic is also useful if we think about the groups as countries—in other words, defining different groups by geography. Using Theil's statistic to measure inequality, we can break global inequality into the proportion of inequality that occurs within countries (i.e. differences in income between US citizens, Chinese citizens, etc.) and inequality that occurs because average income differs between countries (i.e. per capita income is higher in the US than in China). As we will see, this allows us to think about how much of global inequality is attributable to inequality that exists because of unequal economic and social conditions across individuals within countries, and how much global inequality is associated with unequal economic and social conditions in different countries.

IN CONCLUSION

So this is probably the most tedious chapter in the book—sorry to lead off with it—but unless we can understand what kinds of inequality we are talking about and how we measure them, we will never be able to understand the more interesting questions related to current trends in inequality and its impact on our lives.

After reading this chapter, you might come to the conclusion that inequality can't be measured—there are too many different ways to define inequality, too many data problems, and no simple measure that can capture its complexity. But I think that this is the wrong conclusion to draw. There are good reasons to be skeptical about economic data, but not to be cynical. As we are going to see in the next chapter, it is true that different data give us different pictures about the *level* of inequality, but the *changes* in inequality are much more apparent. Whether we are talking about equality as measured by market income, disposable income, wealth, pay, consumption, individuals, households, ethnicity, gender, country, survey data, tax data, citizenship, urban, rural, social mobility, Gini coefficients, or income shares, our confidence intervals regarding the data are shrinking over time, the data is becoming more correlated, and the *trends* in inequality have become pretty clear. Describing these trends in economic inequality is the subject of the next (and much more interesting) chapter.

NOTES

1. US Bureau of the Census (2016).
2. Doepke and Tertilt (2016).
3. Pew Research Center (2014b).
4. Pew Research Center (2011).
5. Hsieh et al. (2018).
6. See Edlund and Kopczuk (2009) and Freund and Oliver (2014).
7. Wolfers (2015).

8. Pager (2003).
9. Watkins (2013).
10. ICP data can be found at: <http://siteresources.worldbank.org/ICPEXT/Resources/ICP-2011-report.pdf>. For a more detailed description of how PPP exchange rates are used to adjust income surveys across countries, see Lakner and Milanovic (2013).
11. You might ask why don't market exchange rates always reflect actual differences in purchasing power? There are three principal reasons. First, market exchange rates are quite volatile and are subject to large swings—20 percent or more in short periods of time—even though the actual domestic purchasing power of the currency has changed very little. The second reason why there is often a difference between market and PPP exchange rates is that many governments manipulate exchange rates and keep them above or below levels to which they would gravitate in the free market. For example, until recently, China's central bank kept the RMB exchange rate undervalued against the dollar in order to facilitate its export markets, meaning that the market exchange rate for the RMB was below its PPP exchange rate. Third, exchange rates over time tend to adjust to reflect the fact that internationally traded goods should sell at roughly the same price across countries. The more goods within a country that are not internationally traded—local foods, real estate, services, etc.—the more likely that a market exchange rate that only equilibrates the prices of internationally traded goods can be consistent with persistently large differences in the price of local goods across countries.
12. Naughton (2017).
13. Galbraith (2012).
14. Atkinson (2015).
15. Hlasny and Verme (2019).
16. *The Economist* (2018a).
17. Korinek et al. (2006).
18. Kopczuk (2015).
19. Sen (1985).
20. Attanasio and Pistaferri (2015).
21. Milanovic (2016).
22. How exactly is a Gini coefficient calculated? Imagine a graph that plots the cumulative population—up to 100 percent of the population—on the vertical axis and cumulative amount of a resource—up to 100 percent—on the horizontal axis. If we plotted what percentage of the population got what percentage of the resource in a world of perfect equality, the distribution would be a 45-degree line and the area under the line would be a triangle. In a world of perfect inequality, there would be no line and no area, only a single point where one person got 100 percent of the resource. The Gini coefficient, then, is a geometric measure of how far away the actual distribution of the resource is from this “ideal” triangle. A higher Gini means that the actual distribution of resources is farther away from perfect equality. See also Peltzman (2009) if you don't like my description of calculating a Gini and would like to read someone else's.
23. World Bank (2018).
24. See Scheidel (2017) for a more detailed description of how these calculations are made.
25. For a brief tutorial on calculating Theil's T statistic, see the University of Texas Inequality Project (2018) at <http://utip.lbj.utexas.edu/tutorials.html>.