1

Demand for Redistribution in the Age of Inequality

If history is any guide, excessive economic inequality never goes down without a fight. Quite literally so: In the past, only mass warfare, a state collapse or catastrophic plagues have significantly altered the distribution of income and wealth (Scheidel, 2018). Could this time be different? With the spread and deepening of democratic institutions, political systems are better equipped today than in the past to reflect the economic interests of the majority of voters and peacefully address, even if imperfectly, high levels of income inequality.

This more optimistic take implicitly assumes that public opinion will act as a countervailing force to rising inequality. For many social scientists, this seems reasonable. As resources concentrate in the hands of a minority, it becomes increasingly advantageous for the poorer majority to redistribute income by taxing the richer minority to fund transfers and public goods (Meltzer and Richard, 1981). As a result, support for income redistribution is expected to increase with income inequality. This increase should be especially large among people at the bottom of the income ladder who have the most to gain from progressive taxation and redistributive spending. Scholars are not alone in expecting the public to react to rising inequality. Pundits and commentators make similar predictions, though, in their case, the motive they impute to voters is rarely economic self-interest. While left-leaning pundits point to voters' moral outrage in the face of "unfair" income differences,[1] right-leaning commentators tie growing support for income redistribution to envy and resentment.[2] Whether due to voters' material self-interest, moral outrage or envy, expectations converge: Greater wealth and income inequality should lead to greater demand for an egalitarian policy response.

[1] "Sorry Washington Post, Bernie Sanders Is Right about Economic Inequality" by John Nichols, in *The Nation*, July 2, 2019.
[2] "Income Inequality and Bullsh*t" by William Irwin, in *Psychology Today*, November 15, 2015.

Still, evidence of rising support for redistribution, especially among the worse off, is scant. As described in more detail in this introduction, the overall pattern is one of striking long–term stability. In the two Western countries with the sharpest increase in income inequality, Great Britain and the United States, any evidence of attitudinal change goes against expectations. In Great Britain, aggregate support for redistribution has not increased but *decreased*. In the United States, attitudinal differences between low-income and high-income voters are *decreasing*, not increasing. How can these contradictory empirical patterns be reconciled with reasonable assumptions regarding the economic determinants of redistributive preferences? What can we conclude regarding public opinion's role as a countervailing force to rising inequality?

This book aims to answer these questions. In Part I, I show that mass attitudes toward redistributive social policies are shaped by at least two motives: material self-interest and fairness reasoning. On the one hand, people support policies that, if implemented, would increase their own expected income. On the other hand, people also support policies that, if implemented, would move the status quo closer to what is prescribed by shared norms of fairness. Combined, these two motives help explain why people often hold redistributive preferences that seem to cut against their own economic interest, with the poor being sometimes opposed to, and the rich very often in favor of, redistributive social policies.

In Part II, I examine how fairness reasoning and material self-interest interact with contextual factors to help explain stability and change in attitudes toward redistributive social policies. I show how, in Western democracies, changes in partisan dynamics have combined with fiscal stress to erode support for key redistributive features of the welfare state. Overall, the evidence suggests that this time might not be so different after all. Without a strong egalitarian turn in mass attitudes toward redistributive policies, there are few reasons to expect the democratic process to bring about ambitious policy responses to rising inequality.

In this introductory chapter, I first present stylized facts regarding expected and observed trends in mass attitudes toward redistributive social policies. I then present the book's main argument and its relationship to the existing literature. I end with a brief description of the chapters to follow.

The Dynamics of Support for Redistribution: Expectations and Evidence

A common expectation is that greater economic inequality will be partially offset by higher demand for policies that redistribute across income groups. In what I will call the "benchmark model," Meltzer and Richard (1981) helpfully

formalize a set of scope conditions and assumptions under which such expectation holds. This model is not designed to capture reality in its complexity. Instead, it provides an internally consistent theoretical benchmark against which to compare and assess the empirical evidence. Any mismatch between the evidence and the model's predictions can be investigated by probing the model further. What does it overlook? How often are scope conditions met? I start with a brief review of this benchmark model and then turn to evidence of attitudinal change in postindustrial democracies.

The Benchmark Model

In the benchmark model, redistributive policies take the form of a flat rate tax and a lump sum per capita transfer equal to total revenue divided by population size. Income inequality is a situation in which some people receive a share of income that is larger than their share of the population ("the rich"), while others receive a share that is smaller ("the poor"). Mechanically, when there is income inequality, the combination of a flat rate tax and a lump-sum transfer results in income redistribution. That's because the tax an individual pays is proportional to their share of national income (high for the rich, low for the poor), while the transfer they receive is proportional to their share of the population (the same for both rich and poor). As a result, the rich pay more in taxes than they receive in transfer. The converse is true for the poor.

A key parameter in this benchmark model is the difference between one's own market income and mean market income, defined as national market income divided by population size. Mathematically, anyone who receives a share of national income that is larger than their share of the population is someone whose own market income is higher than the mean market income. This person will always favor a 0% tax rate as any positive tax rate will result in a net loss, that is, a tax bill that is larger than the transfer received. Conversely, anyone whose market income is lower than the mean market income stands to benefit from a high tax rate. Assuming no administrative costs and disincentive effects, this person will even support a 100% tax rate as the transfer received (equal to mean market income) will always more than compensate for the individual market income lost to taxes.[3] With this redistributive set up, the closer someone is to the bottom of the income ladder, the more they stand to gain. Conversely, the closer someone is to the top, the more they stand to lose.

[3] This assumes no disincentive effects from taxation and no bureaucratic costs. Relaxing these assumptions does not change the intuition presented here.

The comparison between mean market income and *median* income[4] captures whether a *majority* would benefit from a higher tax rate. Indeed, if median market income is lower than mean market income, then a hypothetical 100% tax rate would advantage a majority of the population. If the difference between the median and the mean is large, that is, if a small minority receives the bulk of market income, then not only does a majority stand to benefit from a high tax rate, it stands to benefit a lot. For this majority group, the resulting lump-sum transfer will more than compensate for the higher tax bill. In other words, the number of people who stand to benefit from redistribution and the extent to which they stand to benefit increase with a top-heavy rise in income inequality.[5]

This benchmark model generates two testable predictions. The first one is a positive relationship between the mean-to-median market income ratio and aggregate support for redistribution. The second prediction is a comparatively larger increase in support for redistribution among those closer to the bottom of the income distribution and no increase in support for redistribution among those closer to the top. Importantly, and in accordance with Occam's razor, this model lays out the key institutional and individual-level assumptions (also called micro-foundations) that underpin the expectation of a pro-redistribution turn in countries with rising inequality.[6] These assumptions include a tax and transfer system designed to be redistributive and citizens who prefer more disposable income than less, are informed about rising income inequality and are aware of its implications for their own position as net winners or losers of redistribution. As I show in the following section, when brought to the data, this benchmark model does not perform very well. Building on this evidence, I then revisit some of the model's key assumptions.

Testing the Benchmark Model

The rise in income inequality started in the 1970s, a decade marked by the end of the postwar economic boom and by a crisis of profitability, investment and productivity, as well as stagflation. The policies adopted to address the

[4] Median income is the income of the individual who splits the population into a bottom poorer half and a top richer half.

[5] The concentration at the top pushes the mean income up without affecting the median, thus increasing the gap between the two.

[6] The benefits of engaging with this benchmark model go beyond analytical clarity and tractability. Western societies are built on the ideal of equal dignity, which stands in tension with the existence of income inequality. Given this, a model hypothesizing that democracies have a built-in inequality moderator rooted in voters' selfish pursuit of more equal outcomes is an appealing starting point. It is a way for researchers to join the public conversation without taking a position on the tension between democratic ideals and existing levels of economic inequality.

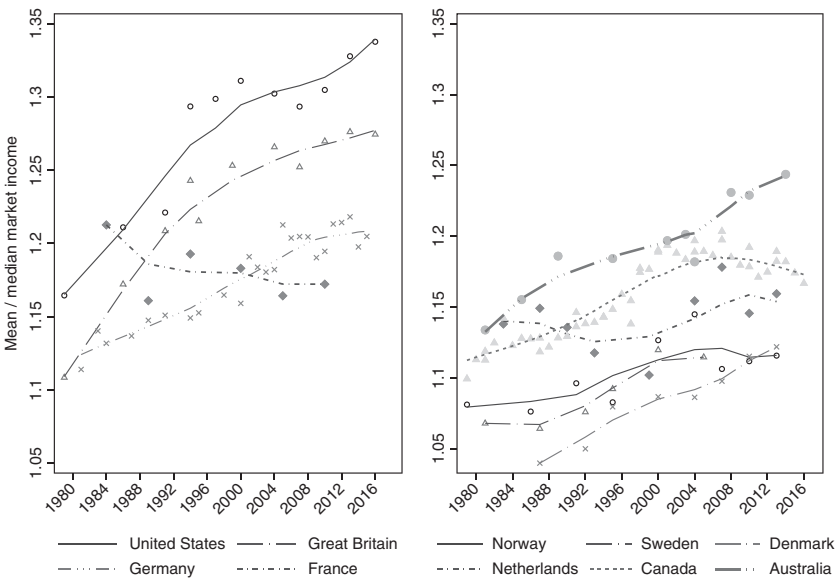

Figure 1.1 Mean-to-median market income
Plots the ratio of mean-to-median gross market income. To improve comparability across countries, the income measure only includes private – not public – pensions. Source: UNU–WIDER, World Income Inequality Database (WIID), www.wider.unu.edu/project/world-income-inequality-database-wii

crisis restored profits and crushed inflation while also contributing to rising economic inequality. These developments have affected some countries more than others. Figure 1.1 plots overtime changes in market income inequality using the mean-to-median income ratio. The figure on the left plots this ratio for all the countries examined with some detail in this book, namely the United States, Great Britain, France and Germany. The figure on the right plots the same ratio for a mix of countries for which similar data are available. The increase in income inequality is most striking in the United States and Great Britain. While positive, the rate of increase in Germany is comparatively lower. France is an outlier: Over the period, the ratio of mean-to-median income is mostly stable (another exception is the Netherlands). Overall, most countries are experiencing an increase in market income inequality.

The mean-to-median ratio obscures what is happening at the two ends of the income distribution. Figure 1.2 plots the average income (market income and public pensions) in the top decile (between the 90th and the 100th percentiles), divided by the average income in the second decile (between the 10th and the 20th percentiles). I focus on the second decile to address concerns that the first

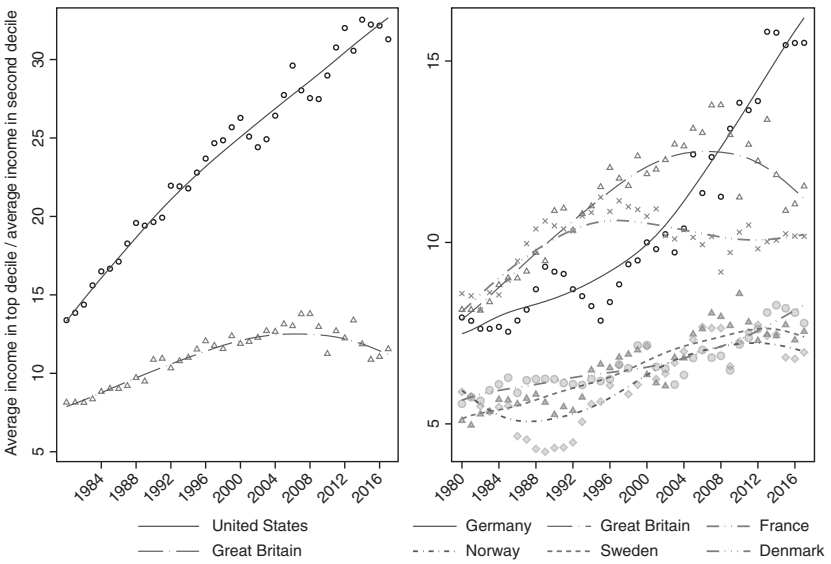

Figure 1.2 Mean market income of the rich relative to that of the poor
Plots the ratio of mean market income in the top decile to mean market income in the second decile. To extend the analysis beyond the working age population, the market income measure includes both public and private pensions. Excluding pensions returns a similar picture, with one exception: the increase in market income inequality in Great Britain is steeper (Atkinson, 2008).
Source: World Inequality Database (WID.world), https://wid.world/data/.

decile might consist of a very disparate group of individuals (e.g., long-term unemployed and students). As shown on the left-hand side, the United States is a clear outlier: Today, the average income in the top decile is thirty times that of the average income in the second decile, representing a tripling of the top-to-bottom income ratio since the early 1980s. In that regard, the evolution in Great Britain is far less dramatic; the average income in the top decile is "only" twelve times that in the second decile, representing a mere 50% increase in the top-to-bottom ratio relative to the 1980s.[7] The figure on the right-hand side plots trends in France, Germany and three Scandinavian countries (Great Britain is included as a benchmark, notice also the change in the *y*-axis). While most countries are experiencing an increase in income inequality, this increase is among the largest in Germany, with France again being the stable outlier.

In light of the trends plotted in Figures 1.1 and 1.2, the United States and Great Britain are ideal candidates for testing the benchmark model. Based on the latter, aggregate support for income redistribution should increase as in-

[7] There is also a noticeable reversal starting with the onset of the Great Recession.

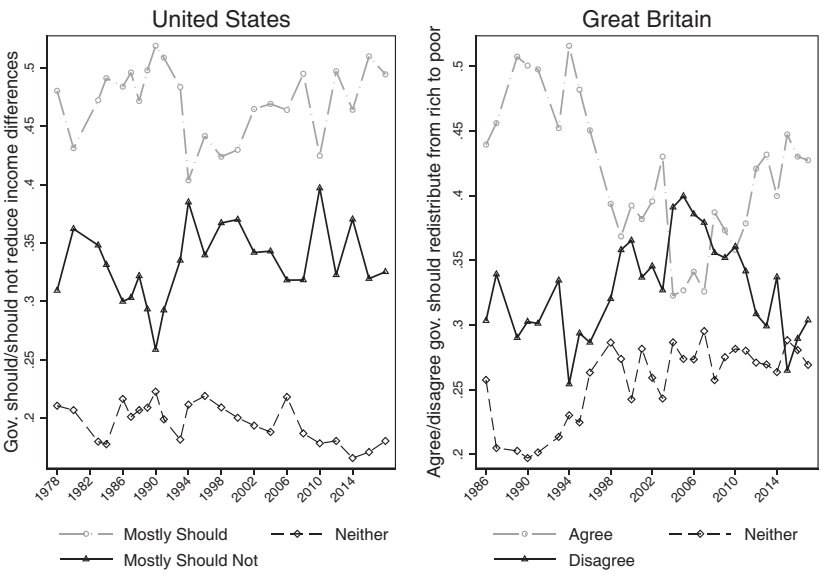

Figure 1.3 Demand for redistribution in the United States and Great Britain
Plots the share of respondents who express mostly support, mostly opposition or neither to a statement asking about income redistribution by the government. Left panel (US): "Some people think that the government ought to reduce the income differences between the rich and the poor, perhaps by raising the taxes of wealthy families or by giving income assistance to the poor (1). Others think the government should not concern itself with reducing this income difference between the rich and the poor (7). (...) What score between 1 and 7 comes closest to the way you feel?" Variable recoded as follows: 1 through 3 "mostly should concern itself," 4 "neither," 5 through 7 "mostly should not." Right panel (GB): "Government should redistribute income from the better off to those who are less well off." Answers recorded using a strongly agree (1)–strongly disagree (5) Likert scale. Variable recoded as follows: 1 and 2 "agree," 3 "neither," 4 and 5 "disagree."
Source: GSS 1972–2018, weighted (left panel); BSAS 1983–2017, weighted (right panel).

come inequality increases, starting with the bottom half of the income distribution. Empirically, this implies an increase in the share of individuals who agree that "the government should redistribute income from the better off to those who are least well-off." Over time, we can also expect attitudinal differences between the top and the bottom of the income distribution to increase. Do we observe the expected increase in mass support for redistribution? Have the preferences of the rich and the poor diverged over time, especially so in the United States?

Overall, the evidence that trends in mass social policy preferences align with theoretical expectations is scant. As shown in Figure 1.3, in the United

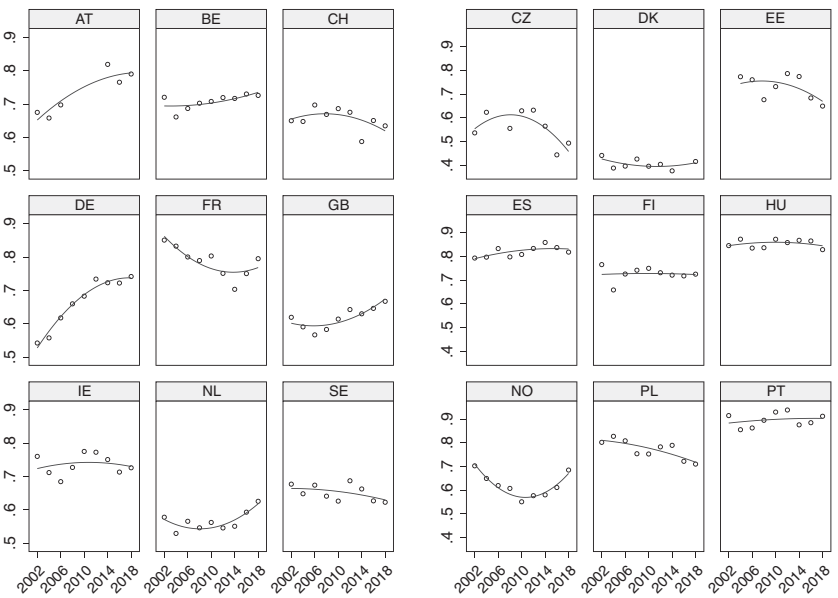

Figure 1.4 Demand for redistribution in postindustrial democracies
Plots the share of respondents who agree with the following statement: "The government should take measures to reduce differences in income levels." Answers were recorded using a strongly agree (1)–strongly disagree (5) Likert scale. In this figure, "strongly agree" and "agree" responses are combined.
Source: ESS 2002–2018, weighted.

States (left panel), the overall pattern is one of striking stability: Despite a sharp growth in income inequality since the 1970s, support for redistribution has remained very stable. In Great Britain (right panel), and against all expectations, the evidence points to a decline in support for redistribution (Georgiadis and Manning, 2012; Grasso et al., 2019). More generally, as shown in Figure 1.4, attitudinal stability is not specific to the United States: In most countries, the trend in support for redistribution is surprisingly flat. One exception is Germany, where support for income redistribution has gone up at the same time as income inequality has increased.

As the rich increasingly stand to lose from redistribution and the poor increasingly stand to win, is there any evidence of diverging attitudinal trends at each end of the income distribution? Figure 1.5 plots the share of respondents in the bottom income quintile who support income redistribution minus the share of respondents in the top quintile who also support it. In both countries, low-income respondents are more likely to support income redistribution than

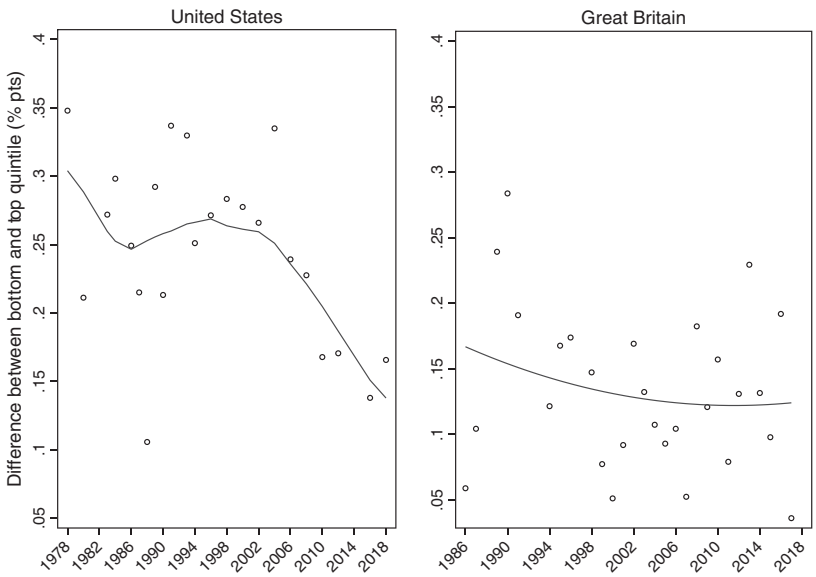

Figure 1.5 Demand for redistribution in the United States and Great Britain: Top versus bottom quintiles
Plots the difference between the share of individuals in the bottom quintile who agree with the policy principle of income redistribution and the share of individuals in the top quintile who also agree. For example, a positive value of 0.2 means that (1) the share of people in the bottom quintile who agree is larger than the share in the top quintile who agree and (2) the difference between the two group shares is equal to 20 percentage points. See Figure 1.3 for item wording. Income measures are described in Appendix A1.1.
US source: GSS, 1972–2018, weighted; GB source: BSAS 1983–2017, weighted.

high-income respondents. In Great Britain, this difference is stable over time. Strikingly, in the United States, the difference between the bottom and the top quintiles is decreasing.

To summarize, despite generational replacement, major recessions, large shifts in unemployment and changing policy paradigms (Hall, Kahler and Lake, 2013), support for redistribution is very stable. In our two most likely cases, Great Britain and the United States, any evidence of attitudinal change goes against common expectations: a decrease in aggregate support in Great Britain and a decrease in the attitudinal income gradient in the United States. Also noteworthy is the difference between Great Britain and Germany, two countries with similar increases in income inequality but with opposite attitudinal trends. Interestingly, France, despite no increase in income inequality, is one of

the few countries (with the possible exception of Spain) to have experienced a nation-wide year-long social movement – *Les Gilets Jaunes* – focusing on economic issues and asking for more income redistribution. How to make sense of these puzzling patterns and country cases? Answering this question requires returning to the benchmark model's micro-foundations: What does the model get wrong, and how can it be amended to get things right?

The Argument Part 1: New Micro-Foundations

In Part I of this book, I relax two of the benchmark model's assumptions. One is the assumption that voters have a sophisticated understanding of their position as net winners or losers of changes to redistributive policy and that it affects their policy preferences. The other is the emphasis, in the form of a fixed rate tax and a lump-sum transfer, on policies' redistributive *consequences*. Relaxing these assumptions suggests a new set of micro-foundations, one in which fairness reasoning takes a leading role.

Fairness Reasoning

The benchmark model's assumption that people are well-informed, self-interested income maximizers is most helpful when economic stakes are quantifiable and large. In countries with mature welfare states, this is rarely the case. First, the redistributive implications of a given policy change are far from straightforward, and politicians, fearing a backlash from affected populations, have only limited incentives to provide clarifying cues. In addition, in countries with mature welfare states, many policy reforms have ambiguous implications (i.e., diffuse costs or benefits), meaning that, for many voters, redistribution is an uncertain or low-stakes issue, with few incentives to acquire the correct information regarding implications for their pocketbooks (Jacobs and Matthews, 2017; Roth, Settele and Wohlfart, 2022). In such a context, the assumption that voters are fully informed selfish income maximizers is heroic at best, requiring researchers to think more creatively about core behavioral motives guiding attitude formation and change.

In this book, I emphasize fairness reasoning as a behavioral motive well suited to the low personal stakes or high-uncertainty world of redistributive politics. Indeed, when it comes to preferences over broad categories of redistributive policies, it is often easier and more rewarding to reason according to fairness principles than to reason based on hypothetical implications for one's own pocketbook. People consequently support policies that move the status