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.

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

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

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

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.

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

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.

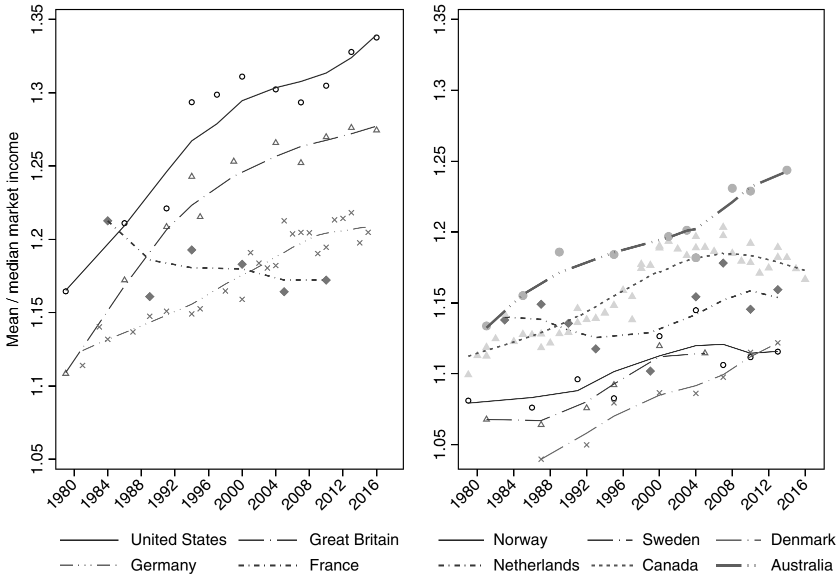

Figure 1.1 Mean-to-median market income
Plots the ratio of mean-to-median gross market income. To improve comparability across countries, the income measure only includes private – not public – pensions. Source: UNU–WIDER, World Income Inequality Database (WIID), www.wider.unu.edu/project/world-income-inequality-database-wii

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

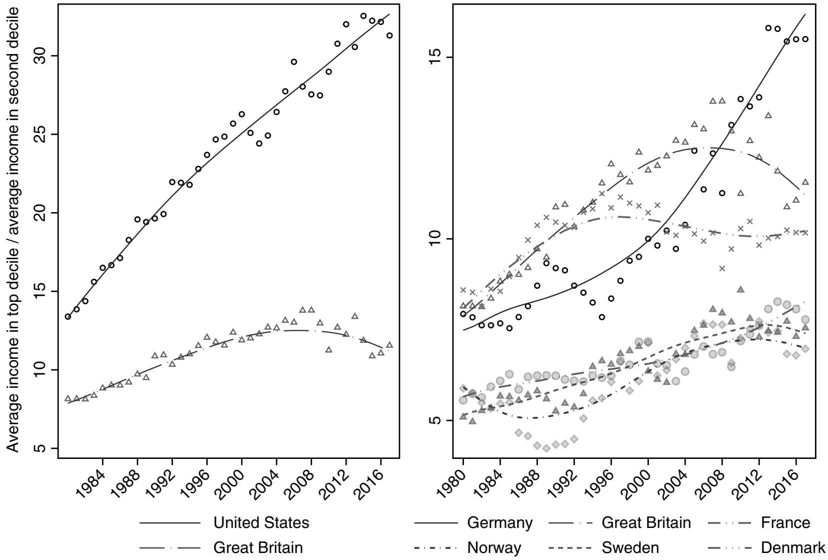

Figure 1.2 Mean market income of the rich relative to that of the poor
Plots the ratio of mean market income in the top decile to mean market income in the second decile. To extend the analysis beyond the working age population, the market income measure includes both public and private pensions. Excluding 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.

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

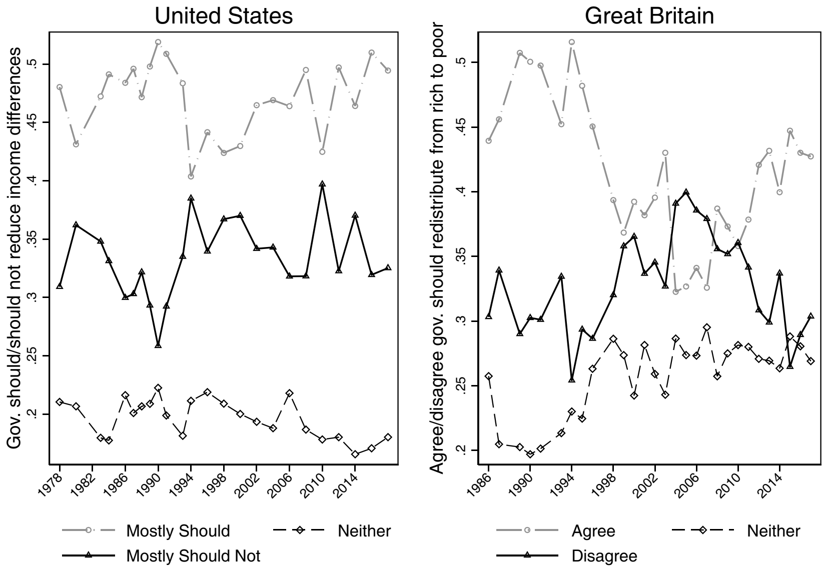

Figure 1.3 Demand for redistribution in the United States and Great Britain
Plots the share of respondents who express mostly support, mostly opposition or neither to a statement asking about income redistribution by the government. Left panel (US): "Some people think that the government ought to reduce 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

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

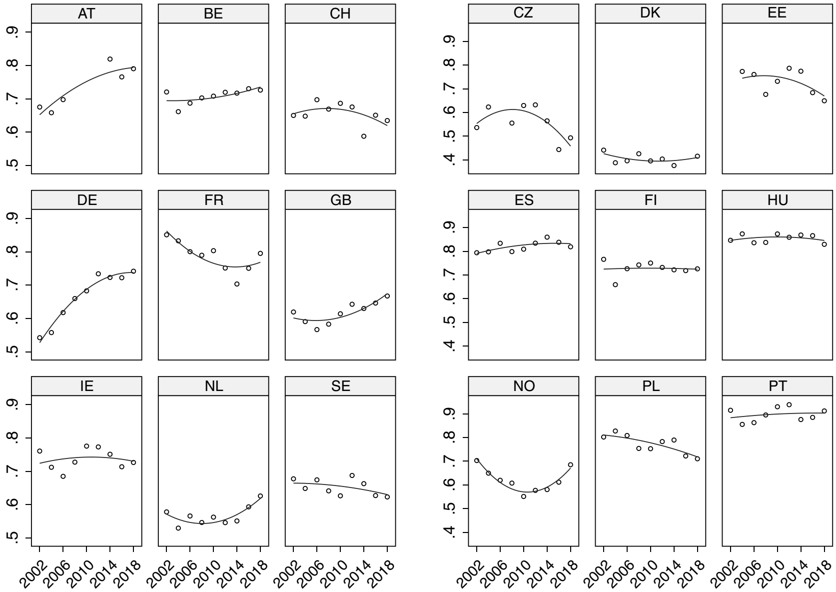

Figure 1.4 Demand for redistribution in postindustrial democracies
Plots the share of respondents who agree with the following statement: "The government should take measures to reduce differences in income levels." Answers were recorded using a strongly agree (1)–strongly disagree (5) Likert scale. 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

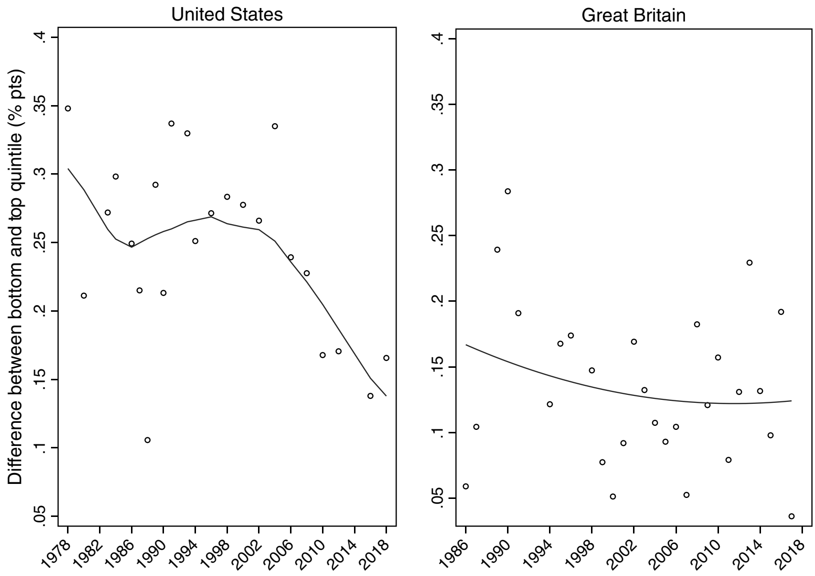

Figure 1.5 Demand for redistribution in the United States and Great Britain: Top versus bottom quintiles

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

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

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 quo closer to what is prescribed by widely shared norms of fairness and oppose policies that move the status quo further away from what is fair. Because people hold different empirical beliefs regarding the fairness of the status quo, they also disagree over which policies to support or oppose. Given that beliefs about the fairness of the status quo are often disconnected from an individual's own position in the income distribution (more on this below), fairness reasoning only incidentally produces the types of policy preferences one might expect from self-interested citizens. Only when policies have large material consequences will people consider what is economically self-serving instead of what is fair. In a world of high uncertainty and low personal stakes, fairness considerations come first and self-regarding considerations about one's own pocketbook come second.

Before theorizing how fairness reasoning and material self-interest combine to shape attitudes toward redistributive social policies, I first provide a more precise definition of fairness reasoning and flesh out its consequences for how researchers conceptualize (and operationalize) attitudes toward redistributive social policies.

Two Norms of Fairness for Two Facets of Redistribution

I define fairness reasoning as the thought process through which individuals act as if a third-party judge ruling on the fairness of a given situation and acting to maximize fairness accordingly. In the context of this book, fairness is maximized by favoring a policy change that moves the status quo closer to what is prescribed by shared norms of fairness. I emphasize two norms of fairness. One is the proportionality norm, which prescribes that individual rewards be proportional to effort and talent. The other is the reciprocity norm, which prescribes that cooperative behavior be rewarded more than uncooperative behavior.

The proportionality norm underpins economic activity in a capitalist society. Consent is achieved when people believe that rewards are proportional to merit, itself a combination of personal decisions as a free agent, individual work ethic, acquired skills and innate talent. The reciprocity norm, in contrast, underpins social solidarity, that is, the provision by the group of basic material security for its members, lest the group not survive temporary material shocks. The successful provision of social solidarity implies group members who are willing contributors to a resource-pooling endeavor and feel no resentment toward those who benefit more than they contribute. This is achieved when people believe that the status quo abides by what the reciprocity norm prescribes, namely the belief that net beneficiaries are cooperators down on their luck, not free riders. In Western democracies, reasoning about the fairness of redistributive

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

social policies consequently implies at least two types of fairness evaluation: (1) How fair is it for some to make (a lot) more money than others in the marketplace, and (2) how fair is it for some to receive more benefits than they pay in taxes? Each question calls to mind a different norm of fairness: the proportionality norm for the former and the reciprocity norm for the latter.

The benchmark model's emphasis on policies' redistributive *consequences*, while correct from an accounting perspective, overlooks this institutional and moral dualism. In doing so, it obscures important distinctions between at least two types of redistributive policies. One type includes policies that directly interfere with high-earning individuals' capacity to generate and keep market income. The other type includes policies that make social insurance redistributive. Policies that fit in the first category include progressive income taxes as well as industrial policies or corporate governance policies that affect the distribution of profits across stakeholders (also called pre-distribution policies). As a shorthand, I call this family of policies *redistribution from* policies because they mostly affect the accumulation of market income by "economic winners." Policies that fit in the second category include social programs that provide a safety net for those who cannot provide for themselves (e.g., means-tested transfers) as well as design features that regulate the extent to which access to generous social insurance programs is conditional on risk profiles and prior contributions (often described as social solidarity). As a shorthand, I call this family of policies *redistribution to* policies because they mostly affect the material conditions of "economic losers." In times of rising income inequality, both *redistribution from* and *redistribution to* policies contribute to more egalitarian outcomes.

The proportionality norm is most prevalent when reasoning about the fairness of *redistribution from* policies. *Redistribution from* policies that interfere with the "natural" allocation of income across economic actors are justified if this allocation is "unfair." In this context, "unfair" implies the belief that market rewards are not proportional to effort and talent. The reciprocity norm is most prevalent when reasoning about the fairness of *redistribution to* policies. *Redistribution to* policies, which redistributes resources to those who cannot provide for themselves (whether temporarily or not), are justified if pooled resources are allocated in a "fair" way. In this case, "fair" means the belief that *redistribution to* policies do not reward free riders more than cooperators.

Support for both types of redistributive policies is the highest among people who find market mechanisms unfair *and* believe that social insurance, especially its most redistributive features, benefit people who are not intentionally trying to "abuse the system." Relatedly, opposition to income redistribution as a whole is highest among people who find market mechanisms fair *and* believe that free riding is highly prevalent. Yet, as I will show in later chapters, most people exhibit a mix of beliefs, making them more likely to hold, for example, liberal preferences toward *redistribution from* policies and conservative preferences toward *redistribution to* policies (or vice versa).

Needless to say, the framework provided in this book should be understood as providing ideal types, not an exhaustive classification tool applicable to all existing norms of fairness, redistributive policies and their possible combinations. Take, for example, unemployment insurance: In most countries, it is designed to be fair according to both the proportionality norm (benefits are tied to contributions and thus to efforts and skills) and the reciprocity norm (the lucky subsidize the unlucky). Individual attitudes toward proposed reforms are consequently shaped by both people's perceptions of how much effort pays in the market place and one's beliefs about the prevalence of free riding. Still, as I will show throughout the book, the simplified framework I propose here has enough traction to justify overlooking some of these nuances.

Why Proportionality Beliefs Need Not Align with Reciprocity Beliefs

I trace the disconnect between proportionality beliefs (beliefs regarding deviations from what the proportionality norm prescribes) and reciprocity beliefs (beliefs regarding deviations from what the reciprocity norm prescribes) to differences between the moral economy of market economies and that of welfare states. Simply put, market economies are morally construed as an *individual* race, while welfare states are constructed as *collective* resource-pooling endeavors among citizens of the same country. One key implication is that, while welfare states imply group boundaries and interdependence, market economies do not. As a result, proportionality and reciprocity beliefs form each according to their own separate rationale.

Proportionality beliefs, I show, are partly self-serving: Individuals who are deemed losers in the individual economic race are more likely to interpret its rules as unfair, while winners are more likely to find it fair. Perceptions of the fairness of resource pooling follow a different logic: They are shaped by differences in how people intuitively reason about social dilemmas, membership and free riding. These differences in how people reason about these issues have little to do with people experiences as winners or losers of the economic race, explaining the disconnect between proportionality and reciprocity beliefs.

Empirically, this means that proportionality beliefs are correlated with one's own labor market position (e.g., income, occupation and skill level). Reciprocity beliefs, in contrast, are correlated with attitudes toward the death penalty, sentencing and discipline in school. These attitudes are commonly used to place people on a libertarian–authoritarian values scale. The disconnect between proportionality and reciprocity beliefs thus partly reflects the fact that liberal–authoritarian values and income are not correlated.

While people's policy preferences are often better predicted by their long-term beliefs about the fairness of the status quo than by a policy's implications for their own pocketbook, material self-interest still matters. First, the correlation between proportionality beliefs on the one hand, and earning capacity on the other, suggests that people form proportionality beliefs in a self-serving fashion, a point I will come back to when discussing how fairness beliefs form and change. Second, while material self-interest does not explain how reciprocity beliefs form, it helps understand *when* and for *whom* reciprocity beliefs matter the most and *when* and for *whom* they matter the least. I explain this latter point next.

Bringing Self-interest Back In

While economic stakes are small and uncertain for many people, for some people and in some contexts, redistributive policies do have large and quantifiable material implications. In such cases, the assumption that individuals behave in ways that maximize their income yields important predictive power. One simple way of conceptualizing the interaction between fairness reasoning and material self-interest is to think of attitude formation as following a two-step reasoning. First, fairness beliefs help people decide whether to support or oppose policies that seek to change the status quo. Second, material self-interest helps predict how much individuals deviate from what fairness reasoning prescribes. As a result, the policy preferences people ultimately settle on are a function of the type of fairness belief they start from on the one hand, and the policy's effect on their own disposable income on the other.

This is particularly relevant for *redistribution to* preferences. Remember that large portions of the morally conservative poor believe that *redistribution to* policies benefit undeserving recipients and that the prevalence of free riding is high. Because more low-income individuals benefit from *redistribution to* policies, they are also more likely, relative to high-income individuals, to adjust their support in a self-serving fashion. As a result, in surveys, the rich appear more driven by fairness concerns than the poor, who are much more likely to give the self-interested answer to a question about *redistribution to* policies.

A lot has been said about the irrational poor who vote against their interests. The argument and framework presented in the book show that this statement does not easily extend to redistributive preferences. On the one hand, low-income respondents are more likely to find the status quo unfair due to beliefs about the high prevalence of free riding. On the other hand, they are less likely to translate these beliefs into support for less generous *redistribution to* policies. High-income respondents, in contrast, are more likely to be morally liberal and to believe that *redistribution to* policies benefit deserving recipients. They are also more likely to translate these beliefs into support for generous and inclusive *redistribution to* policies. As a result, attitudes expressed by this group are rarely in line with their "objective" material self-interest, especially when compared to the attitudes of low-income respondents.

The Argument Part 2: Explaining Attitudinal Stability and Change

The goal of this book is to explain patterns of attitudinal stability and change in times of rising income inequality. Part II builds on the conceptual framework presented in Part I to theorize mechanisms of mass attitudinal change and document their role in the country cases discussed at the beginning of this chapter.

Based on this framework, attitudinal change results from the activation of one (or more) of three possible causal pathways.[8] One pathway is a change in the share of the population who, when asked about redistributive social policies, give the self-interested answer instead of the fair one (or vice versa). Such a change in the dominant mode of reasoning is most consequential when it affects cross-pressured individuals: those whose fairness-maximizing response differs substantively from their self-interested one. A second pathway is a change in the fairness rule people rely on when answering a given set of survey questions. This causal mechanism is best known in the literature under the umbrella expressions of *framing* or *priming* effect. It is most relevant when a given issue can be framed in ways that activate either proportionality concerns or reciprocity concerns. Because of differences in proportionality and reciprocity beliefs, answers can change depending on which norm is activated. A third pathway is a change in fairness beliefs. Unlike framing effects, which tend to only have temporary implications, a change in fairness beliefs has lasting and substantive consequences.

[8] As a friendly reminder, in this book, attitudinal change is a change in patterns of answers to *survey questions* about redistributive social policies, questions that have been asked repeatedly over time in large national and cross-national surveys.

Theorizing aggregate attitudinal change requires understanding the contextual factors that activate each of these three causal pathways. In this book, I discuss three activating factors. The first is fiscal stress, which affects attitudes through a reversal to self-interest among altruistic net contributors to *redistribution to* policies. The second consists of survey design and item wording, which affect attitudes through framing effects. The third is party competition dynamics, which affect attitudes by enabling changes in fairness beliefs. In Part II, I demonstrate how these activating factors help explain the puzzling patterns highlighted at the beginning of this chapter.

Fiscal Stress and the Erosion of Social Solidarity in France

In most countries, majority support for generous *redistribution to* policies is achieved through a combination of self-interested support among beneficiaries and altruistic support among those who, while not benefiting, nevertheless support it because it is the "fair thing to do." Fiscal stress, I argue, can trigger a reversion to self-interest among these altruistic supporters, leading to a decrease in support. As the threat of fiscal adjustment becomes more credible, people consider the implications for their own pocketbooks and increasingly favor letting the generosity of targeted policies mechanically erode. Absent fiscal stress, and because of fairness reasoning,[9] risk exposure and income are poor predictors of attitudes toward *redistribution to* policies. However, in times of fiscal stress, reversion to self-interest implies the re-emergence of a relationship between risk and income on the one hand, and support for more generous *redistribution to* policies on the other.

To test this argument, I focus on a most likely case, France, where tax increases are a credible threat. Since the 1990s, tax increases have occurred at regular intervals, always to deal with fiscal imbalances tied to social insurance.[10] Thanks to the availability of a rich decades-long longitudinal survey, I can examine the relation between fiscal stress generated by the Great Recession, and support for *redistribution to* policies, specifically willingness to pay for generous social transfers targeted to the poor and support for making access to social insurance unconditional on prior contributions.

The French case not only offers an ideal setting for testing a causal pathway, it also speaks to my argument's explanatory power. Given proportionality beliefs extremely critical of markets and inequality (documented in Part I), *Les*

9 Indeed, a large subset of high-income and low-risk individuals trust that beneficiaries of *redistribution to* policies are "deserving" recipients.

10 In contrast, in the United States, the threat of tax increases has remained elusive due to legislative gridlock and the country's ability to borrow at very low interest rates.

Gilets Jaunes movement's emphasis on more *redistribution from* is to be expected. Yet, as documented in Part II, these beliefs do not protect against the eroding effects of fiscal stress on support for *redistribution to* policies. The French case thus further highlights the need to distinguish between the two facets of demand for redistribution.

Great Britain: A Decrease in *Redistribution to*, not *from*

Policy debates about income redistribution are inherently multifaceted: They can emphasize fairly or unfairly taking from the rich, or fairly or unfairly giving to the poor. Empirically, this dualism has implications for how researchers interpret changes in answers to a survey item that asks about redistribution from the rich and redistribution to the poor in the same breath. This type of survey item is commonly found in longitudinal surveys and has become ubiquitous in studies of attitudinal change. In the book, I call this survey item the "traditional redistribution item." Because beliefs that make it fair to redistribute from the rich are distinct from beliefs that make it fair to redistribute to the poor, answers to this survey item can change depending on contextual primes and framing effects, making survey answers tricky to interpret.

I show that contextual primes tied to survey design partly explain the puzzling decline in the share of British respondents who agree that the "government should redistribute income from the better off to those who are less well off." In the British Social Attitudes Survey (BSAS), used in Figure 1.3, the redistribution item is included right after a block of survey items asking about free riding among the unemployed and the poor. Because reciprocity beliefs in Great Britain have much less egalitarian implications than proportionality beliefs, respondents primed to think about redistribution as a *redistribution to* policy express lower levels of support than when such prime is absent. As a result, the decline in support for redistribution documented in Figure 1.3 has to be interpreted with caution: It reflects a conservative shift in reciprocity beliefs (more on this below) and, consequently, a decline in support for *redistribution to* policies specifically, not *redistribution from* policies. As with the French case, these findings further highlight the pitfalls of conceiving of demand for redistribution in the singular.

Framing effects matter beyond survey design. Indeed, decades of research have shown that partisan competition, through priming and framing effects, shapes how people reason about a given redistributive issue (Zaller, 1992). Assuming a mid-to-long-run shift in partisan competition, it can also directly impact people's fairness beliefs, the third causal pathway examined in this book.

https://doi.org/10.1017/9781009366038.001 Published online by Cambridge University Press

Partisan Dynamics and Belief Change in the United States and Great Britain

By fairness *beliefs*, I mean the subset of claims and considerations people hold that jointly shape their assessment of the status quo as barely deviating, deviating somewhat or deviating a lot from what the proportionality or reciprocity norms prescribe. For example, in the case of the proportionality norm, this includes the belief that "effort pays," that the education system is "meritocratic" or that "hiring is based on skills not personal connections." For some people, these beliefs might be only partially consistent, suggesting that these individuals do not have a clear sense of the fairness of the status quo. Yet, as I show in Chapter 3, most people answer survey items in ways that allow them to be ranked from holding "consistently unfair" considerations to "mostly unfair," to "mostly fair," to "consistently fair" considerations. To understand how people's "basket" of considerations change, I build on Zaller (1992) and argue that belief change happens at the intersection of (1) changes in one's discursive context, that is, the claims and considerations one is exposed to, and (2) one's propensity to accept and incorporate new claims into one's own basket or resist doing so despite exposure to new claims.

To identify systemic changes in people's discursive context, I focus on mid-to-long-term changes in party competition that affect how competing political parties talk about redistributive issues. People mirror these changes in predictable ways depending on their own exposure to elite messaging and partisan affinities. Furthermore, I show that people experiencing hardship are more likely to resist dissonant right-wing considerations about the fairness of market institutions and inequality and embrace consonant left-wing ones. The same logic applies to reciprocity beliefs: People who are moral liberals are more likely, relative to people who are moral authoritarians, to resist dissonant right-wing considerations about the prevalence of free riding and accept consonant left-wing ones regarding the deservingness of recipients. I trace these differences to cultural differences in how people reason about social dilemmas. These mechanisms of exposure and resistance help explain why proportionality beliefs correlate with income and why reciprocity beliefs correlate with liberal–authoritarian values.

This simple model of belief formation sheds new light on changes in mass attitudes on both sides of the Atlantic. In Great Britain, disagreement over how to better foster economic growth generated important tensions between the traditional left and the pro-market "Third Way," with the latter winning control of the Labour party in 1995. The first-past-the-post system meant that the traditional left was eclipsed by Tony Blair's winning coalition. The consequence of the demise of the traditional left in favor of the New Labour was a depoliticization of proportionality concerns in favor of a center–right bi-partisan consensus on the need to reform *redistribution to* policies. I document how mass attitudinal trends in Great Britain have mirrored these supply-side dynamics, that is, relative stability with regard to proportionality beliefs and a widely shared conservative shift in reciprocity beliefs. As reciprocity beliefs have become more conservative, support for cuts to welfare transfers has increased. In line with my argument, the impact on support for *redistribution to* policies is the largest among high-income respondents; low-income respondents, in contrast, maintain comparatively lower and stable support for cuts.

I contrast these trends with those found in Germany, where, after the Social Democratic Party (SPD's) shift to the right under Gerhard Schröder, the traditional left was, thanks to proportional representation, able to maintain its political relevance, with implications for the discursive context. This helps explain why in Germany attitudinal trends following Schröder's electoral victory have looked very different from that found in Great Britain after Blair's victory.

While Great Britain's main parties experienced a convergence on redistributive issues, in the United States, parties have famously polarized. In the United States, underpinning overall attitudinal stability is a pattern of mass partisan polarization that mirrors what is happening at the elite level. This is particularly true of reciprocity concerns, with the largest implications for the attitudes of high-income Democrats. I document a pro-redistribution shift in patterns of answers among high-income Democrats, a group increasingly likely to find minorities,[11] and low-income individuals in general, deserving recipients of *redistribution to* policies. Low-income Republicans, in contrast, have been much less likely to translate growing concerns over free riding and minorities' access to benefits into opposition to *redistribution to* policies. As a result, and against common expectations, the decline in the income gradient originates in growing support for redistribution among rich Democrats, not declining support among poor (often white) Republicans.

[11] It is by now well documented that in the United States, beliefs about the fairness of *redistribution to* policies are highly racialized: People disagree over the extent to which minorities, especially African Americans, are deserving beneficiaries of collectively pooled resources. Students of American politics have studied these beliefs under the umbrella term of symbolic racism (also called racial resentment). As shown by Enders and Scott (2019), and in line with the argument presented in this book, symbolic racism exists separately from adherence to liberal-conservative principles about government intervention and the fairness of market forces. See Chapter 8 for more detail.

Full Circle

In light of dramatic increases in income inequality, why aren't voters asking for more income redistribution? One explanation is fairness beliefs' stabilizing role as an *anchoring* proto-ideology in a context in which stakes are low or uncertain. Absent triggering factors, attitudinal stability is the norm. Rising inequality does not constitute a triggering factor for at least two reasons. First, for most people, income inequality is an abstract reality, meaningful only through the lenses of fairness reasoning and prior fairness beliefs. Only a subset of the population – the one that already finds market income unfair and already support redistribution – will experience a rise in income inequality as something needing to be addressed. Another reason is the existence of more than one type of fairness evaluation: In the realm of *redistribution to* policies in particular, there is no obvious theoretical or empirical connection between market income inequality and policy preferences.

A strong egalitarian turn in mass attitudes toward redistributive policies will require a pro-redistribution change in fairness beliefs, an increase in the share of people who support such policies out of material self-interest or both. Under what conditions might this happen? Will these conditions be met any time soon? Answers to these questions partly depend on political parties' positioning on redistributive issues as well as the state of the government's finances.

My argument suggests that party competition dynamics have provided a discursive environment inhospitable to a pro-redistribution turn. In the 1980s, in response to stagflation, a new consensus emerged built on the assumption that increasing the reach of market mechanisms, both geographically and in terms of what counts as tradable goods and services, could help improve paltry growth rates. Increasing the reach of market mechanisms does not imply deregulation. Instead, it requires a state-driven process of institution building, the paragon being the European single market. In the process of crafting more markets (Vogel, 2018), elites have affected the discursive context in ways unfavorable to egalitarian fairness beliefs, specifically proportionality beliefs. Indeed, over the past few decades, disagreement over the merits of markets (and the fairness of market income inequality) has become less common, decreasing the likelihood of mass changes in the belief that the status quo deviates from what the proportionality norm prescribes.

Contrast this relative depolitization of *redistribution from* policies to what has happened with *redistribution to* policies: In both Great Britain and the United States, the redistributive features of the welfare state have been the target of extensive rhetorical and policy innovation ("from welfare to workfare," "flexicurity," welfare chauvinism, etc.). As a result, aggregate attitudinal change on the *redistribution to* dimension has mirrored patterns of elite convergence (in the case of Great Britain) and polarization (in the case of the United States). In the concluding chapter, I tentatively argue that the restructuring of electoral cleavages (from being centered on income and class to being centered on education and skills) has made it electorally less risky to debate the need for more or less *redistribution to* than to disagree over the need for more or less *redistribution from*.

In Great Britain, the drastic austerity measures taken to address fiscal stress have likely accelerated the erosion of support for *redistribution to* policies. In the United States, the threat of tax increases has remained elusive due to legislative gridlock and the country's ability to borrow at very low interest rates. This helps explains the importance of fairness reasoning among high-income Americans.

Based on the argument presented in this book, the Covid crisis opens contradictory possibilities. On the one hand, there is little doubt that beneficiaries of *redistribution to* policies are not opportunistic free riders. Yet large levels of public spending have likely triggered concerns over moral hazard. The resulting debt, if it becomes politicized, could erode nascent support for generous and inclusive social transfers and programs among Democrats. In the realm of pre-distribution and taxation policies, stock market rallies and successes of companies such as Amazon raise sharp fairness concerns that echo those heard after World War II against war profiteers (Scheve and Stasavage, 2012). Yet absent entrepreneurial politicians willing to play this populist left-wing card, we are unlikely to witness any sharp change in how people reason about *redistribution from* policies.

Existing Literature and Competing Theories

I am not the first person to emphasize the importance of fairness reasoning. There is, for example, a long line of work examining the role "beliefs in a just world" or perceptions of the origins of poverty (Alesina and La Ferrara, 2002) play in generating support for the status quo (Benabou and Tirole, 2006; Lerner, 1980). Still, these studies cannot account for important stylized facts, most importantly differences between support for *redistribution from* policies and support for *redistribution to* policies. In line with research by Hvidberg, Kreiner and Stantcheva (2020), I find that people most likely to win in the economic "race" are also more likely to find it fair. What this research overlooks is that these same economic winners are also central to redistributive coalitions that support *redistribution to* policies because they believe that this is the fair thing to do.

The part of the argument regarding the mediating role of material self-interest is closely related to work by Rueda and Stegmueller (2019) who argue that, because of declining marginal return to consumption, only the rich have the luxury to reason altruistically. I extend this reasoning in at least two ways. First, instead of altruism, I emphasize the role of fairness reasoning, which might or might not result in altruistic preferences. Indeed, for a subset of the population, retrenchment – and not expansion – is the "fair" policy. Second, I emphasize the role of uncertainty over the personal consequences of redistributive social policies. My argument generates new predictions, unaccounted for by Rueda and Stegmueller, such as why, for some policies (e.g., Earned Income Tax Credits), it is the rich who are divided based on what they believe to be fair, while for other policies (e.g., a wealth tax), it is the poor.

How does my argument compare with previous attempts at explaining the missing left turn? According to one set of contributions, there is no reason to expect higher inequality to translate into higher levels of support for redistribution because voters simply do not behave like rational actors seeking to maximize their own income (Bartels, 2005; Sears and Funk, 1991). In contrast, I show that reality is more nuanced and that material self-interest plays a key role. Other studies have examined how contextual factors, such as residential segregation, shape people's (mis-)perceptions of inequality and their own position in the income distribution (Cruces, Perez-Truglia and Tetaz, 2013; Sands, 2017; Sands and de Kadt, 2020). Such studies assume a very specific distribution of fairness beliefs, one in which the median voter finds the status quo unfair, as defined by the proportionality norm. Only then can one expect correcting people's perceptions of inequality to translate into growing demand for more egalitarian policies. Whether or not a majority perceives income inequality as a violation of the proportionality norm is something to be explained, not assumed.

Studies that connect lower or declining support for redistribution to the presence of poor minority groups suggest an alternative interpretation of the empirical patterns described in this book (Alesina and Glaeser, 2006). According to this line of work, the decline in support for redistribution documented in Great Britain follows from growing hostility toward immigration, which famously culminated in the Brexit vote. From this perspective, differences in reciprocity beliefs are mostly capturing differences in people's hostility to minorities. Throughout the book, I show that this line of reasoning under-delivers both empirically and theoretically. Empirically, a necessary condition for net contributors to withdraw their support for policies that redistribute to the poor is that they perceive immigrants as overrepresented among net beneficiaries. However, there is little correlation between the latter and demand for more or less *redistribution to*. One reason is that the perception that the modal recipient is an immigrant (shared by many) does not imply the belief that immigrants do not deserve to benefit from *redistribution to* policies (most common among moral conservatives). Beyond tautologically arguing that some people are more xenophobic than others, few contributions provide useful tools to unpack why people differ on the latter point. In contrast, the interpretative framework presented in this book helps explain why moral conservative are both less likely to find immigrants deserving and more concerned about moral hazard and free riding in general. In doing so, it unpacks an important overlap between "second dimension" attitudes and redistributive preferences overlooked by existing work. This overlap sheds a new light on the evolution of far-right nativist movements, one that helps unpack the fairness concerns central to these movements' success.

Outline

The rest of this book explains my argument in greater detail and provides empirical support for my claims. In Chapter 2, I define fairness reasoning, explaining what it is and why it matters. In Chapter 3, I flesh out and test fairness reasoning's implications for how researchers conceptualize and measure demand for redistribution broadly speaking. In Chapters 4 and 5, I bring material self-interest back into the picture. Chapter 4 shows that, while proportionality beliefs appear to be self-serving, reciprocity beliefs do not and instead overlap with attitudes and values associated with "second-dimension" preferences. Chapter 5 documents how fairness reasoning and material self-interest combine in predictable ways depending on individual characteristics and policy design. This concludes the book's first part.

Jointly, the chapters in Part I propose a new interpretative framework, selected first by combining knowledge from across the social sciences and second by examining how well this framework fits the data relative to existing interpretative frameworks. The strength of the demonstration hangs on the number of empirical facts better accounted for by one theory over the other.

Part II focuses on explaining patters of stability and change both at the individual and country levels. In Chapter 6, I lay out the framework's implications for changes in mass attitudes toward redistributive policies. I discuss three causal pathways and three related activating factors. One activating factor is fiscal stress, which affects attitudes through a change in the mix of motives. It is discussed in Chapter 7 using the case of France. Another is survey design, which affect observable attitudes through framing effects. The third factor is party competition dynamics, which can affect mass attitudes by changing fairness beliefs. Both are discussed in Chapter 8 with the case of Great Britain and the United States. Chapters 9 and 10 zoom in on individual-level belief change, documenting the role of material self-interest for proportionality beliefs and moral worldviews for reciprocity beliefs.

Empirically, I test the existence of hypothesized individual-level pathways using survey experiments and regression models that leverage within-individual variation. At the country level, I strive to provide plausible accounts of changes in mass social policy preferences in the United States, Great Britain, Germany and France. Plausibility is assessed by leveraging longitudinal data (causal arguments have clear implications for the sequencing of events), qualitative accounts of important critical junctures (e.g., a change in elite discourse following Tony Blair's election) and by addressing alternative explanations (e.g., the role of immigration, discussed in Chapter 6).[12]

Chapter 11 concludes by discussing this book's implications for understanding the public's heterogeneous response to rising income inequality.

Appendix

A1.1 Measuring Income

Great Britain: Measuring Income in the BSAS

BSAS respondents were asked to provide an assessment of household income from all sources by choosing among a list of income brackets. New top income brackets were regularly added throughout the years. First, I transform the income intervals into their common-currency midpoints (e.g., 2,000–3,000 becomes 2,500). Second, for the top category, I use the method recommended by Hout (2004), which imputes an income value that is a function of the number of respondents in the top category and the number of respondents in the bracket that precedes it. This information, combined with a few assumptions regarding the skew of the income distribution, seeks to compensate for underestimating income levels among those with the highest income in the sample. Finally, I multiply the estimate by the GDP deflator available on the Bank of England website and adjust for household size by dividing the resulting amount by the square root of the number of people living in the household. I then divide this

12 The book draws on both small and large-N methodological traditions in political science. For a discussion of how these methodologies combine, see the online appendix available at www.charlottecavaille.com.

income measure into year-specific quintiles. Income trends in the BSAS are similar to those found in Labour Force Surveys: In the BSAS, the gap between the mean bottom quintile household income and the mean top quintile household income has increased from a multiple of 9 in 1986 to a multiple of 15 in 2009.

United States: Measuring Income in the GSS (and ANES)

Given the much sharper rise in inequality in the United States, I need a different empirical strategy than the one recommended by Hout (2004). Following McCarty, Poole and Rosenthal (2008), I use data on the actual distribution of household income to estimate the expected income within each categories in ways that better match the shape of the income distribution in a given year. Indeed, for some income categories, the midpoint might either over- or underestimate the average income of individuals in this category relative to the shape of the income distribution. The computation of these estimates is detailed in McCarty, Poole and Rosenthal (2008) (see the appendix in Chapter 3), and final individual measures were obtained through the program Matlab. This method for approximating income extrapolating from a categorical income variable assumes that the actual income is distributed following a log–normal distribution with a time-varying mean and variance obtained using labor force survey data. After imputing a continuous income measure for each respondent, I multiply this value by the GDP deflator available on the Federal Bank's website. I then adjust for household size following the US Census Bureau's recommendation to assign different weights to adults and children and to allow weights to decrease with each additional adult or child. I divide the resulting income measure into year-specific quintiles. As a robustness check, I also run the analysis using the year specific P20 and P80 measures provided by the Census Bureau to identify people in the top and bottom quintiles.