How the Geographic Clustering of Young and Highly Educated Voters Undermines Redistributive Politics[*]

Tom O'Grady[†]
University College London

Andreas Wiedemann[‡]
Princeton University

September 4, 2023

Short title: Geographic Clustering and Redistributive Politics

Abstract

We analyze support for the welfare state across time and space in Great Britain. Using multi-level regression and post-stratification with historical data and an original survey, we show that a virtually identical majority of people supported those policies in the mid-1990s and in 2020, but patterns of support were very different. Young and highly-educated people are now the strongest supporters, as are the youngest and most highly-educated geographic areas, mirroring divides over 'second-dimension' issues like Brexit. However, young and highly-educated voters are clustered in a small number of places, with the Labour party struggling to win moderately-educated and moderately-young areas. As a result, the Left's problem in majoritarian systems is not the rise of second-dimension politics per se, but rather how its supporters are distributed spatially along that dimension. A *majority of voters* in favor of welfare and redistribution no longer translates as easily into winning a *majority of places* in support.

Keywords: political geography; redistribution; education

Many countries are experiencing a growing divide in economic fortunes between thriving urban centers of the knowledge economy and peripheral "left-behind" places. These rifts have political implications. The Brexit referendum and the election of Donald Trump represented the "revenge of places that don't matter" (Rodríguez-Pose 2018), fueled by declining social mobility, fading job prospects, and feelings of political neglect (Baccini and Weymouth 2021; Cramer 2016; Gimpel et al. 2020). In this paper, we argue that in the UK—a typical majoritarian electoral democracy— socio-economic changes have also made it harder for left parties to enact social policies. The spatial clustering of younger and highly-educated people, who are now the new core supporters of welfare and redistribution, undermines redistributive politics when combined with a plurality voting rule in single-member districts.

Recent work has documented how the spatial concentration of voters matters for policy preferences (Beramendi 2012; Enos 2017; Rodden 2010; Wiedemann 2023) and political representation (Döring and Manow 2017; Jusko 2017; Rodden 2019). Prior work on support for social policy and redistribution, though, has focused either on nation-wide changes or variation across regions. A lack of suitable data has prevented studies of geographic variation in preferences at a more localized scale and over time. As a result, we know very little about the spatial distribution of preferences toward welfare and redistribution, how this has shifted over time, and what this means for the ability of parties to enact redistribution. We fill these gaps with new evidence from Great Britain, drawing on historical survey data from the 1994-6 waves of the British Social Attitudes Survey (BSAS) and an original nationally representative survey from 2020. This allows us to study fine-grained changes in the geographic distribution of preferences over a 25-year period. We take as given that the geographic clustering of social groups affects the average preference of an area through the direct effects of sorting by social groups, as well as peer effects and contextual effects. Our focus is on measuring the political consequences of these processes: the average preference for welfare and redistribution in every local authority area in Great Britain in both periods, using multilevel modeling and post-stratification (MRP) (Hanretty, Lauderdale and Vivyan 2018; Warshaw and Rodden 2012). In addition to measuring combined support, we also disaggregate preferences into a redistribution dimension and a deservingness dimension.

We use our estimates to demonstrate the impact of three key socio-political developments on the politics of redistribution. First, educational upgrading means that there are many more *uni-*

versalists in the population: people who are highly-educated, younger, and socially liberal (Born-schier et al. 2021). As we show in this paper, these groups have now also become key supporters of welfare and redistribution. By contrast, twenty-five years ago, low-income and unemployed people were much more numerous and formed the core of support. Second, this means that the ability of center-left parties to form a geographic coalition in favor of redistribution is crucially dependent—in majoritarian countries—on where younger and highly-educated people live. We demonstrate that they are unevenly distributed in space. A few places like cities and university towns have many very young and highly-educated residents, but many other places have far fewer. As a result, support for welfare and redistribution is, from the perspective of left parties, much more inefficiently distributed today than it was when low-income and unemployed people were the main supporters of redistribution twenty-five years ago. Third, political competition in many cases now takes place along a libertarian-authoritarian axis—so-called "second-dimension" politics—over issues such as immigration that do not necessarily align with the "first-dimension" of conflict over economic issues and the size of government (Caughey, O'Grady and Warshaw 2019; Häusermann and Kriesi 2015; Inglehart 1990). In countries with majoritarian electoral regimes, the challenge in assembling a geographic majority for redistribution is to win the median area in terms of support. But these median areas today have stronger support for such second-dimension issues that are a poor fit with center-left policies. Today it is more difficult for the left to win these median-supporting areas than in the past, when political conflict revolved more around economic issues.

This combination of changing individual-level social policy preferences, geographic clustering, and a new axis of political competition undermines the translation of a majority of preferences in favor of the welfare state into legislative seats and political power. The political challenge for the left is not so much the rise of second-dimension politics per se; we find that nation-wide popular support for redistribution was largely unchanged from the mid-1990s to 2020. It is rather the way that its voters are distributed spatially in terms of that dimension. Thus in a country like Great Britain, where elections are held in single-member districts under plurality rule, a *majority of voters* in favor of welfare and redistribution no longer translates as easily into winning a *majority of places* in support. The rise of the knowledge economy, the changing landscape of economic opportunity, and the clustering of socially-liberal young and highly-educated people (Boix 2019; Iversen and

Soskice 2019) have important implications for redistributive politics by shaping electoral coalitions when elections are held under plurality rule.

Our findings also speak to current debates about growing regional inequalities and potential policy solutions for so-called "left-behind" areas. Because support for welfare and redistribution is now so strongly tied to social conservatism, and because left-behind places are relatively socially conservative, places that would benefit the most from government redistribution are the least supportive of such policies, even though the residents of highly successful and wealthy places might be willing to pay for it. In fact, we find no evidence for any geographic realignment of party support in terms of areas' support for welfare and redistribution in the UK. The Conservative party continues to perform best in areas with the weakest support for redistribution, implying no greater electoral pressure for them to redistribute. This helps explain why the current Conservative government has proposed to "level-up" economically declining areas through investment in infrastructure, jobs, and education rather than focusing on traditional transfer-based redistributive policies (Jennings, McKay and Stoker 2021). Meanwhile, in geographic terms, Labour continues to be the party of redistribution—thanks to its new heartlands in cities and university towns where young and highly-educated people are clustered. As the knowledge economy further amplifies the economic divides between successful urban areas and the remaining periphery, the political dynamics we document in this paper are likely to become more important in the future.

Age, Education, and Support for Welfare and Redistribution

The transition from Fordist manufacturing to the knowledge economy has marginalized many regions, notably former industrial heartlands. Economic decline, fading job opportunities, and political resentment against urban elites are part of the reason why the votes for Brexit in 2016 and Donald Trump in 2020 were concentrated in non-urban areas and, more broadly, why populism has thrived in "left-behind" places (Adler and Ansell 2020; Baccini and Weymouth 2021; Cramer 2016; Gimpel et al. 2020; McKee 2008; Rodríguez-Pose 2018). On the other hand, knowledge-intensive industries, together with adjacent service-sectors, have agglomerated in urban centers and university towns. Well-educated and young people have clustered in those places, leaving "left-behind" areas—often former industrial regions—relatively old and bereft of graduates (Ansell and Gin-

grich 2021; Boix 2019; Iversen and Soskice 2019). Our aim is to understand how this clustering of young and well-educated people in space affects the distribution of support for redistribution and welfare across small geographic areas. Before introducing our geographic preference estimates, in this section we explain the connections between age, education, and policy support. The clustering of young and well-educated people shapes places' average demand for social policies through geographical sorting as well as peer and contextual effects, which we discuss in turn.

The Impact of Geographical Sorting by Young and Highly-Educated People

Direct effects of sorting refer to the impact of, for instance, university graduates moving to a new area and maintaining their old opinions after the move. The impact of these dynamics on the geography of preferences for redistribution and welfare depends on the preferences of different age and education groups. Throughout this paper, we estimate peoples' preferences for redistribution and welfare based on the analysis of five questions that were originally asked in the British Social Attitudes survey (BSAS). The BSAS is the UK's gold standard for a nationally-representative social survey, which has been asking these questions almost annually since 1985. All five questions have 5-point Likert response scales. The specific question wordings are as follows *(original item abbreviations from the BSAS in parentheses)*, asking "how much do you agree or disagree that ...":

- If welfare benefits weren't so generous, people would learn to stand on their own two feet? *(welffeet)*
- The government should redistribute income from the better-off to those who are less well off? *(redistrb)*
- The government should spend more money on welfare benefits for the poor, even if it leads to higher taxes? *(morewelf)*
- Many people who get social security don't really deserve any help? *(sochelp)*
- Most people on the dole[1] are fiddling in one way or another? *(dolefidl)*

The panels in Figure 1 show over-time patterns for each question by education group[2] and by three age groups—young, middle-aged, and elderly—from 1985 to 2019/2020 (the latest available data).[3] Degree-holders have been more likely to perceive welfare recipients as deserving or

---

[1]"Social security" and "the dole" in the UK generally denote benefits and transfers and are widely-used terms, equivalent to the American use of "welfare."

[2]'Secondary school and below' includes people who received qualifications at the age of 16-18 (O-level, GCSE, A-level and equivalents) or lower, including those with no qualifications.

[3]We present education data only up to 2019, since a change in the question wording about education in 2020 renders 2020 incomparable to the past data.

Figure 1: Public Opinion on Welfare and Redistribution, 1985-2020

(a) By Education

(b) By Age

Notes: The figures show mean responses by groups on 5-point likert scales, where higher values indicate greater support for each statement. *Source:* British Social Attitudes Survey 1985-2020.

better-behaved since the 1980s (*welffeet*, *sochelp*, and *dolefidl*). School-leavers and degree-holdrrs have generally shown about equal support for redistribution and welfare spending (*redistrb* and *morewelf*), although since the mid-2010s degree-holders have moved slightly ahead of the other groups. By 2019 degree-holders were the most supportive group across all five questions, with a

much larger gap for issues related to deservingness. Differences by age have mostly been quite small since the late 2000s, but recently a gap has also opened up between the young and old on all measures except *morewelf*, with young people now more supportive. Thus both young and highly-educated people have become core supporters of redistribution and welfare in the UK.[4]

These subgroup differences are readily explainable using theories of support for redistribution and welfare. To understand why, we must first distinguish redistribution *to* the poor and redistribution *from* the rich (Cavaillé and Trump 2015). Highly-educated people tend to behave more altruistically (Rueda 2018). They hold socially liberal values that make them favor redistributing to the poor, whom they perceive as more deserving, hard-working, and less prone to fraudulent behavior than low-educated people (Attewell 2021; Bullock 2021; Gelepithis and Giani 2020). They are more tolerant of diversity and supportive of immigration, whereas lower-educated natives fear that immigrants will be the target of redistribution: the phenomenon of welfare chauvinism (Häusermann and Kriesi 2015; Hjorth 2016; Mewes and Mau 2013; Reeskens and van Oorschot 2012). On the other hand, concerns about moral deviance from welfare claimants are strongly linked to authoritarian values, which are much more prevalent among lower-educated people (McArthur 2022). Hence beliefs about the deservingness of welfare claimants are more related to positions on the "second" non-economic libertarian-authoritarian dimension of political conflict than to the "first" dimension of traditional left-right conflict about economic issues. In addition, last-place aversion affects low-educated groups close to, but not at, the bottom of the economic ladder. They are motivated to defend their relative privilege by excluding the very poorest from redistribution, with concerns about loss of relative status leading them to support radical right parties that campaign to limit benefits to "deserving" natives (Gidron and Hall 2017; Kurer 2020;

---

[4]An obvious question is whether recent changes by age are temporary or due to methodological issues related to the Covid-19 pandemic. The 2020 BSAS was collected from October to December 2020 using online and telephone surveys, in contrast to previous years when interviews were conducted face-to-face. To help verify whether this affected results, a panel study was carried out in July 2020 to re-contact people who had been interviewed face-to-face in past years (BSAS 2021). Conducted online, this survey revealed near-identical results to the late 2020 final survey, meaning that the "pattern of support [for redistribution] increasing most among those who are younger was in evidence in our July 2020 panel as well as in our most recent BSA survey – and thus is less likely to be an artefact of the change in our survey methodology" (BSAS 2021, p. 24). In most cases the changes also did not begin in 2020 but rather continued trends since the early or mid-2010s. The subgroup patterns were also replicated almost perfectly in our own separate survey that was conducted via YouGov in September 2020, shown later in this paper. The 2021 BSAS data are not yet available for analysis, although a report detailing headline results is available and shows little or no change in aggregate support for the variables that we include (BSAS 2021). Therefore we do not believe that the 2020 patterns were temporary or anomalous, due to the pandemic.

Kuziemko et al. 2014). For all of these reasons, the education gap in the UK for deservingness perceptions—*welffeet*, *sochelp*, and *dolefidl*—holds across European countries (Attewell 2021).

Highly-educated people could also in theory be more opposed to redistributing *from* the rich to the poor, as captured by the *redistrb* and *morewelf* questions. Education leads to higher income, lower unemployment risk, a superior ability to self-insure through savings, and better prospects regarding future job security and income growth. Material self-interest could therefore reduce support for redistribution among highly-educated people (Attewell 2021; Bullock 2021; Gelepithis and Giani 2020; Marshall 2016, 2019). In addition, university education inculcates a set of shared economic norms and expectations that could lead graduates to oppose redistribution (Gelepithis and Giani 2020; Mendelberg, McCabe and Thal 2017). From this perspective, views on "redistribution from" are much more related to the classic left-right dimension of economic issues than to the libertarian-authoritarian axis of conflict. Consistent with these expectations, recent empirical research on European countries from the European Social Survey up to 2016 has found that highly-educated Europeans are more skeptical of redistribution from the rich than their less-educated counterparts (Attewell 2021; Gelepithis and Giani 2020).

Figure 2: Differences in Support for Redistribution between Degree-Holders and School-Leavers

Notes: Differences in mean responses by country and year on 5-point likert scales, where higher values indicate greater support for redistribution. For each country, both years are shown where available; otherwise the latest data from 2019 or 2020 is shown. *Sources:* 2019 data from ISSP Inequality Survey; 2020 data from ISSP Environment survey.

However, although the knowledge economy has benefited those with higher education overall, many highly-educated people today—particularly young people, women, and ethnic minorities—are becoming labor market 'outsiders.' They are employed in the service sector with higher levels of unemployment risk, lower incomes and more insecure employment contracts, with involuntary part-time work also common. This group with high education but low income and economic security has grown rapidly in knowledge-economy societies. It now constitutes the core constituency of left parties and strongly supports redistribution from the rich, together with very left-wing stances on the deservingness of recipients (Abou-Chadi and Hix 2021; Häusermann, Kurer and Schwander 2015; Kitschelt and Rehm 2022). The growing size of this group should lead to a narrowing or reversal of the overall gap between highly-educated and lower-educated people on redistribution, as well as a growing convergence between views on the redistribution and deservingness dimensions. The Covid-19 pandemic has likely strengthened this dynamic because highly-educated people were more exposed to economic risks than in a typical economic downturn. All of this helps explain the recent developments in the UK shown in Figure 1a.

Recent cross-national survey evidence, collected several years after the data from Attewell (2021) and Gelepithis and Giani (2020), supports this view. Figure 2 shows data from the International Social Survey Program (ISSP): a question with different wording to the BSAS (agreement with the statement "it is the responsibility of the government to reduce the differences in income between people with high incomes and those with low incomes"). This question featured in both the 2019 ISSP Inequality Survey and the 2020 ISSP Environment Survey. A substantial number of global north countries are available in the 2019 study; a smaller subset are also available in 2020, with fieldwork occurring during the Covid-19 pandemic in all cases. Today, there is little or no difference in opinions on "redistribution from" between degree holders and school leavers in most places. Only 6 out of the 17 countries show a clear statistically-significant negative difference, with the UK similar to other countries. In several European cases (Denmark, Finland, Germany, and Switzerland) there was a substantial narrowing of the gap during the pandemic from 2019-20. Both the UK-specific and cross-national data suggest that higher support for "redistribution from" among the highly-educated have emerged since 2015 and even more strongly in 2020, so that views on this policy dimension are becoming more aligned with views on deservingness, where

highly-educated people have always been much more left-wing. Today, therefore, the highly-educated now form a core demographic group in support of welfare provision and redistribution.

Young people, particularly in the UK, are also theoretically predicted to be more supportive of welfare and redistribution. They are very socially liberal and, as a result, more likely to perceive welfare claimants as deserving. But they are also much more economically insecure than older generations, who often have generous and secure pension income as well as housing assets. Young people have lower income, less wealth, and weaker access to housing and job security, which makes them more supportive of redistribution from the rich (Bell and Gardiner 2019; Green and de Geus 2021; Pearce and Chrisp 2021). Together with recent changes in support by education, this is resulting in a cleavage in rich democracies that increasingly pits highly-educated, middle-class, socially liberal, and younger "universalists" who favor redistribution and believe that welfare claimants are deserving against low-educated, working-class, older, and more authoritarian "particularists" who believe the opposite (Bornschier et al. 2021). The evidence presented here suggests that in the UK in particular, such patterns have strengthened since the 2010s and particularly during the Covid-19 pandemic. This means that the more that young and highly-educated people sort into a geographic area, the more supportive that area will become of welfare policies.

The Impact of Contextual Effects

The spatial clustering of people with similar education levels or ages—and, therefore, particular views on redistribution and welfare—should further shape preferences through contextual effects. Being surrounded by like-minded people directly influences support for welfare provision through peer effects; or there may be contextual factors of "left-behind" places and knowledge economy hubs, such as their racial diversity or levels of poverty, that shape these preferences too. The existence of compositional effects versus context has been widely debated and measured in the political science literature (Abrams and Fiorina 2012; Cho, Gimpel and Hui 2013; Enos 2017; Gallego et al. 2016; Maxwell 2019; Sands 2017). We do not seek to distinguish the relative importance of contextual and selection mechanisms as drivers of preferences. We take it as given that both are likely to matter. Instead, we measure the spatial distribution of preferences that has resulted from these processes, document changes over time, and explore its political consequences.

Geographic and Political Implications

Both sorting and contextual effects imply that support for redistribution and welfare is more heav-ily concentrated in urban areas and university towns populated by young and highly-educated people, compared to the recent past. These new geographic patterns have strong implications for redistributive politics under majoritarian electoral rules. Prior work has shown that single-member districts bias the translation of votes into seats in favor of right-wing parties as left-wing parties win the few urban districts with very high vote margins while losing many non-urban districts with small margins (Rodden 2010, 2019; Wiedemann 2023). Our argument builds on this research but points to important and understudied spatial changes in the political coalition behind redistributive policies. As education and age have become key predictors of welfare support, the growing spatial clustering of younger and highly-educated voters constrains legislative majorities for redistribution under majoritarian electoral rules. This makes it harder for left parties in majori-tarian countries to provide generous welfare policies, even when nation-wide support for these policies is as high as in the past. In this paper, we provide empirical evidence for these claims.

Measuring Support for Redistribution in Small Geographic Areas

Despite our strong theoretical expectations that the geography of support for welfare and redistri-bution has changed, no evidence currently exists about local-level social policy preferences. This is largely due to lack of adequate data. Measuring public opinion toward welfare and redistribution over time and across space faces two major challenges. First, we need survey questions that mea-sure support consistently over time. Second, we need geo-coded survey data with large enough samples to measure support within political units such as electoral districts or local governments.

Data and Survey Questions

To address these challenges, we combine two sets of survey data. Our data on historical public opinion come from the 1994, 1995, and 1996 waves of the BSAS. Most data on the characteristics of local authority districts is not available until the 1990s, and the BSAS ceased providing the geographic location of its respondents in the late 1990s. The 1994-6 period therefore represents a "sweet spot" that maximizes historical data at the local level while mostly preceding the large

changes in higher educational attainment that have occurred in recent decades. We combine these three BSAS survey waves into one large survey, treating 1994-6 as a single period in order to enlarge the number of observations per local authority. There would be too few respondents in any one of the surveys for reliable estimation using multilevel regression and post-stratification (MRP) but they provide enough observations when combined. Altogether this yields data on 8,797 British adults. Our measure of contemporary social policy preferences comes from an original survey we fielded in September 2020 using the British survey firm YouGov. Unlike current editions of the BSAS, our survey included indicators of respondents' local authority. We obtained a nationally representative sample of 6,314 British adults, providing enough responses to carry out MRP in the aftermath of recent changes in social policy preferences by age and education.

We estimate opinions on redistribution and welfare based on the five original BSAS questions shown above. In our YouGov survey we re-asked these questions with identical wording and response scales. Since both the 2020 YouGov survey and the 1994-6 BSAS surveys are random samples of the British population, we can meaningfully compare change over time between the two. Nonetheless they were collected with different methodologies and by different survey companies. To check whether these methodological and company effects may be important, we compare our 2020 YouGov survey to the 2020 BSAS survey in Appendix Figures A.1 to A.3. The latter was also collected in the fall of 2020, after the outbreak of the Covid-19 pandemic. In almost all cases we find only negligible differences between both surveys. Mean support on our three constructed scales (see below for details) and on the five individual survey questions was virtually identical in both surveys, as were subgroup differences by age and education. In both surveys, young people and highly educated people were more supportive of welfare and redistribution by almost identical amounts. Hence, methodological differences between YouGov and the BSAS do not drive our results or prevent us making comparisons between the two over time.

Prior to our analysis, where necessary we re-coded the variables such that higher values indicated more left-wing opinions (i.e., pro-welfare or pro-redistribution). After these transformations, all five survey items are positively correlated with each other in the historical BSAS data and in our own contemporary data, as shown in Figure 3. We are interested in the two dimensions of redistribution described by Cavaillé and Trump (2015). Redistribution "from" is captured by *redistrb* and *morewelf*, whereas redistribution "to" is captured by *welffeet*, *dolefidl*, and *sochelp*. In

Figure 3: Inter-Item Correlations Among the Survey Questions

BSAS, 1994–6

	redistrb	morewelf	dolefidl	sochelp	welffeet
redistrb	1.00	0.40	0.11	0.16	0.21
morewelf	0.40	1.00	0.26	0.28	0.37
dolefidl	0.11	0.26	1.00	0.58	0.61
sochelp	0.16	0.28	0.58	1.00	0.61
welffeet	0.21	0.37	0.61	0.61	1.00

YouGov, 2020

	redistrb	morewelf	dolefidl	sochelp	welffeet
redistrb	1.00	0.53	0.33	0.36	0.45
morewelf	0.53	1.00	0.48	0.47	0.62
dolefidl	0.33	0.48	1.00	0.64	0.67
sochelp	0.36	0.47	0.64	1.00	0.67
welffeet	0.45	0.62	0.67	0.67	1.00

Notes: Inter-item correlations among survey questions in the 1994-6 BSAS and our original survey of 2020.

both periods the two sets of variables mostly show stronger correlations among themselves than with the other set, suggesting that they do capture distinct concepts. However, in line with the evidence presented in the previous section, the two-dimensional structure is much less evident in 2020, with stronger correlations between all five variables. Today, people who score highly on "redistribution to" are much more likely than in the past to score highly on "redistribution from." We therefore created three dependent variables: (i) "redistribution," which is the average response to *redistrb*, and *morewelf*; (ii) "deservingness," which is the average response to *welffeet*, *dolefidl*, and *sochelp*; and "combined support," which takes the average of all five variables.[5] Comparing the 2020 YouGov and 1994-6 BSAS surveys, we find that mean scores on the combined dimension (3.21 in 1994-6 vs. 3.25 in 2020), the deservingness dimension (3.13 vs. 3.25) and the redistribution dimension (3.32 vs. 3.30) were almost identical in the two periods, as the panels in Figure 1 also suggested. But as we will show below, these extremely similar levels of *aggregate* support translate into very different geographic patterns of support in both periods.

Figure 4 shows how differences in average support for redistribution, deservingness, and our combined welfare measure across groups have changed over time. Low-income people were more supportive than high-income people in both periods on our combined measure, but this masks dif-

---

[5]We use averages rather than, for example, principal component analysis (PCA), to ensure that estimates in both periods are fully comparable. PCA could lead to differences that arise from different variable loadings rather than changes in opinion per se.

Figure 4: Support for Redistribution and Welfare by Groups, 1994-6 vs. 2020

Notes: Average response values with 95% confidence intervals. Higher values indicate greater support. Low and high income are defined as respondents in the bottom and top quintiles of household income, respectively. Secondary education includes all respondents whose highest qualification was completing secondary school (up to A-Level) or lower; university degree includes all respondents with at least an undergraduate degree.

ferences between the deservingness and traditional redistribution dimensions. While low-income people are considerably more supportive of redistribution than high income people, both income groups share similar levels of support on the deservingness dimension. Those with low incomes score lower on deservingness than on traditional redistribution because their greater social conservatism and authoritarianism conflicts with their greater economic egalitarianism (Häusermann and Kriesi 2015; Kitschelt and Rehm 2022). Conversely, due to their distinctive views on deservingness highly-educated people were overall the most supportive group welfare and redistribution in both periods, joined by young people in 2020. Between-group differences were also larger for age and education than for income. By 2020 young and highly-educated people were the core supporters of welfare and redistribution, equaled by low-income people only for the redistribution scale. Overall, today's supporters of welfare and redistribution are very likely to be young and well-educated, especially in a political environment where deservingness is more salient than in the past, in part due to an increased focus on the (perceived) moral failings of welfare recipients in media and political discourse (O'Grady 2022).

Unit of Analysis: Local Authority Districts

We use our two surveys to measure support for welfare and redistribution over time and across space for British local authority districts, which are the lowest level of government. Each district covers a distinct and coherent geographic area such as a city or county.[6] In 2020 there were 368 local authorities in England, Scotland, and Wales (Great Britain); our data does not include Northern Ireland. Most of these areas have populations between 100,000 and 400,000 although there are a few outliers such as the Isles of Scilly (just over 2,000) and Birmingham (over 1.5m). Patterns of party support across local authorities are very highly correlated with patterns by electoral constituency for the House of Commons. We focus on local authorities rather than constituencies for two reasons. First, the technique we use to measure local opinions (MRP) requires samples larger than what we could feasibly collect in order to properly measure constituency-level opinion (Hanretty, Lauderdale and Vivyan 2018), since there are many more constituencies than there are local authorities. Second, we compare geographic patterns of public opinion in the mid-1990s to 2020. Over that period there have been very large boundary changes to electoral constituencies that make it nearly impossible to compare like with like over time.

This is not the case for local authorities. Since the 1990s, changes to existing local authority boundaries have been very modest, to the extent that they can be ignored in our analysis.[7] There have been some mergers of District Councils into Unitary Authorities, reflecting efficiency drives from the central government. This means that there were more local authorities in existence in the 1990s compared to today. Mergers took place at various points, especially in 1996 and again in 2009. Importantly for our purposes, however, these took place in a many-to-one fashion without impinging on existing boundaries. For example, in North-East England a new Unitary Authority called Northumberland Council was created from the previously-existing Alnwick, Berwick-Upon-Tweed, Blyth, Castle Morpeth, Tynedale and Wansbeck District Councils. We developed a lookup table to create our measures with either the districts as they existed in the 1990s or the districts as they existed in 2020. When measuring local authorities with 1990s boundaries we use the

---

[6]For historical reasons, local authorities are variously known as District Councils, Metropolitan Districts, Boroughs (in London), and Unitary Authorities. In some English areas an additional tier of government sits above them, known as County Councils, which do not feature in our analysis. In others a Unitary Authority covers a whole county, which therefore *do* feature in our analysis.

[7]A typical change involves moving one street into a different local authority.

boundaries that existed in the year 1998, since some of our local-authority-level covariates were collected with these boundaries. This means that we measure opinion in 408 local authorities in the mid-1990s, compared to 366 in 2020.[8] Thus with the 2020 boundaries we have a fully balanced panel dataset of opinions over the two periods across Great Britain.

Estimating Support for Redistribution and Welfare at the Local Level

We estimated the average opinion in each local authority using multilevel regression and post-stratification (MRP). The first stage—multilevel regression—involved estimating Bayesian multilevel models. We followed standard approaches, including recent applications to the UK (Hanretty 2020; Hanretty, Lauderdale and Vivyan 2018; Kastellec, Lax and Phillips 2010). For each of our three outcome scales (redistribution, deservingness, and combined support) we estimate the following model for individual i, in local authority j and region m, with level l of each categorical variable, in both periods across n local authority areas:

$$Opinion_{ij} = \alpha_0 + \alpha^{sex}_{l_{sex}[i]} + \alpha^{age}_{l_{age}[i]} + \alpha_{j[i]} + \epsilon_{ij}$$

with:

$$\alpha^{sex}_{l_{sex}[i]} \sim N(0, \sigma^2_{sex}), \text{ for } l_{sex} = 1, 2$$

$$\alpha^{age}_{l_{age}[i]} \sim N(0, \sigma^2_{age}), \text{ for } l_{age} = 1, ..., 9$$

$$\alpha_{j[i]} \sim N(Z'_j\delta + \alpha_{m[i]}, \sigma^2_{la}), \text{ for } j = 1, ..., n$$

$$\alpha_{m[i]} \sim N(0, \sigma^2_{region}), \text{ for } m = 1, ..., 11$$

where $Opinion_{ij}$ refers to the opinion of each individual for one of our three scales in a given time period; α indicates a modeled effect with α_0 as the overall intercept; and Z is a set of local-authority-level predictors with associated coefficients δ. Thus each individual's opinion is a function of their individual demographic characteristics (age and sex) as well as the local authority where they live.

To allow for post-stratification, the individual characteristics match available data on local authorities' demographic composition. We were able to obtain data from 1996 and 2019 on the proportion of each local authority's adult population by sex and age group (18-19, 20-29, 30-39,

---

[8]For example, the score for "Northumberland" in 1994-6 is an average of the measures for Northumberland's predecessor authorities Alnwick, Berwick-Upon-Tweed, Blyth, Castle Morpeth, Tynedale, and Wansbeck. Although there are 368 local authorities in 2020, we produce estimates for 366 because we lack covariate data for the City of London and Isles of Scilly.

40-49, 50-59, 60-69, 70-79, and 80+). Post-stratification precludes us from using further individual-level predictors because population data are not available at a granular enough level. For example, there is no data on local authority population shares defined by the intersection of age, sex, *and* education. Nonetheless, both age and sex are strong predictors of social policy attitudes, and in their methodological study of how to measure opinion in British parliamentary constituencies, Hanretty, Lauderdale and Vivyan (2018) find that the inclusion of strong district-level predictors matters much more for the accuracy of predictions than the post-stratification stage.

Our local authority-level predictors all capture important influences on opinion, including variables that should affect the demand for welfare as well as an area's conservatism. They include an area's benefit claimant share (a proxy for unemployment), share of population that is aged 18-24, share of population that is aged 65+, mean gross disposable household income in pounds, mean amount of benefits received per person in pounds, population density measured as people per square kilometer, percent of the adult population educated to degree level or higher, and mean real house price in pounds. Finally, we obtained the annual change in gross value added in percent (a proxy for economic growth), but only at the level of eleven broad regions that are used by British statistical agencies.[9] The regions also have their own modelled effects, α_m. The uneven availability of data required us to be pragmatic about which dates to use our geographic predictors for; we used the nearest available date in all cases. Appendix Section B details all data sources, including the years for which we use them. Both of the local authority income variables (household income and benefits received) as well as the mean house price were logged. All continuous variables were re-scaled to have a mean of zero and standard deviation of one prior to analysis.

Estimation was carried out using the R package brms (Bürkner 2018) with default non-informative priors. We ran four chains for 2,000 iterations each, with the first 1,000 as warmup. All standard diagnostics indicated excellent convergence. Our results are therefore derived from 4,000 samples of the posterior distribution, using posterior mean predictions for our main estimates. For each opinion scale in both periods, we then estimated each area's average opinion using post-stratification. This involved first predicting the opinion of each demographic sex-age group in each local authority area, and then weighting these by each group's share in the local authority

---

[9]The regions are East of England, East Midlands, London, North-East England, North-West England, Scotland, South-East England, South-West England, West Midlands, Yorkshire & The Humber, and Wales.

district's adult population. We also estimated the uncertainty of our post-stratified local authority estimates by re-running the prediction and post-stratification steps for all 4,000 posterior samples. We emphasize that the choice of Bayesian as opposed to frequentist estimation was due to the ease of accurately modeling uncertainty in the Bayesian setup. We show in Appendix Figures C.1 and C.2 that our results are virtually indistinguishable from an approach that uses frequentist non-Bayesian multilevel models.

Geographic Patterns of Support for Redistribution and Welfare

Using this setup, we created estimates for 408 local authorities over the 1994-6 period and 366 in 2020. We normalize our MRP estimates to the 0-1 interval in all figures. Figure 5 plots the densities of our three measures of local authority-level social policy preferences. Across local authorities, the distribution of opinions on redistribution and welfare has changed very little over time. On average, areas were slightly more supportive overall in 2020 (the left panel) and even more likely to perceive welfare recipients as deserving. There was no change on the redistribution scale over time. In all cases a slight right skew is evident; a few places are much more supportive than the average, but there are fewer outliers at the opposed end of the spectrum. To show the uncertainty of these results, Appendix Figures D.1 and D.2 show the mean estimates by local authority, alongside 95% credible intervals derived by re-running the prediction and post-stratification steps on each posterior sample. The mean predicted levels of support are tightly estimated, especially in 2020.[10] Local authorities in the center of the distribution are, unsurprisingly, not statistically distinguishable from each other but this is not true as we move towards the tails.

Two sets of findings emerge so far. First, popular support for welfare and redistribution was virtually identical in 1994-6 and 2020, despite changing support among education and age groups. Second, the geographic distributions of support were also almost unchanged (Figure 5). However, this apparent geographic stability masks considerable volatility, both between and within areas, since the mid-1990s. The maps in Figure 6 provide a more detailed picture of the geographical distribution of welfare support in Great Britain in the mid-1990s and 2020, using our combined

---

[10]The estimates for 1994-6 are more uncertain in some cases because the survey data excludes around one-third of areas, whose predictions are naturally more variable. In 2020 the survey data contains at least one respondent from every local authority

Figure 5: Distribution of Social Policy Support Across Local Authority Districts, 1994-6 and 2020

Notes: The dashed line indicates the district average for the 2020 YouGov survey; the solid line indicates the district average for the 1994-6 BSAS survey. Estimates are normalized to the 0-1 interval. Higher values indicate greater support.

measure of support. In the 1994-6 period, polarization was mainly between the more conservative midlands and the South of England and the more liberal areas of Wales, Scotland, Inner London and the North of England. The (then) Labour party heartlands in rural or semi-rural North-East England, central Scotland, and South Wales were relatively pro-redistribution. In 2020, however, the highest support was much more likely to be found in the most urban areas. Cities like Cambridge, Cardiff, Edinburgh, Glasgow, Manchester, and Newcastle stand out more from their rural hinterlands than in the past. More outlying areas of London have converged towards inner London too.[11]

Socio-Economic Context and Welfare Support

In the rest of the paper we show why this reshuffling of social policy preferences across areas makes it harder for left parties to win a majority of areas favoring welfare and redistribution, even though a similar number of both people and places supported them. Geographic support

---

[11]Figure E.1 shows separate maps for each of all three redistribution preference dimensions for the mid 1990s and 2020. It shows that the geographic distributions of the three measures of support for welfare became more correlated in 2020. In the earlier period, some areas like South Wales and North East England were more supportive of redistribution than deservingness, but in 2020 the same areas stood out on both dimensions.

Figure 6: Predicted Support for Welfare and Redistribution by Local Authority, Great Britain

1994–6   2020

Notes: Predicted local authority estimates based on 2020 survey data and 1994-6 BSAS data for the combined social policy estimates. Values are normalized to the 0-1 interval. Higher values indicate more support for social policies. Boundaries differed in 1994-6 and 2020.

has realigned along age and education lines. Winning areas close to the geographic median is much more challenging than it was in the 1990s because those areas are moderately-aged and educated—terrain where victory for the left is more challenging as politics has realigned around non-economic issues.

Changes in the Socio-Economic Characteristic of Local Authorities

We begin by examining socio-economic changes across areas. The density plots in Figure 7 show how the socio-economic characteristics of local authorities have changed over time according to five main indicators that also feature in our regression models. We use 2020 geographic bound-

aries and the nearest-available data to 1994-6 and 2020, respectively, as described in the Appendix Section B. Local authority districts have become only marginally more polarized in terms of household income, with all areas becoming richer over time on average. Unemployment rates (measured by claimant shares) have fallen in, and converged across, all local authorities such that there is very little polarization today. In 1994-6, a substantial number of local authorities had extremely high levels of unemployment; the average unemployment rate was also higher across Great Britain. By the late 2020s the distribution was much tighter. No area had very high unemployment, with the peak rate falling from almost 25% in the 1990s to under 10% in 2020.

Meanwhile, the distribution of areas by youth (share of 18-24 year-olds) became substantially more right-skewed between the mid-1990s and the present. Areas were slightly less "young" by 2020—with the average share of 18-24 year-olds falling from 8.2% in 1996 to 7.9%. However, a key difference is that many more areas were very youthful. In 1996, only two local authorities, the university towns of Cambridge and Oxford, comprised more than 15% 18-24 year-olds. In 2020, that number had increased to ten local authorities. Young people have become somewhat more clustered in a smaller number of areas. Average educational attainment in local authorities rose from 16.6% of adults with a degree in the mid-1990s to 42.6% in 2020, reflecting a huge increase in the number of degree-holders in the adult population. These degree holders have continued to cluster in knowledge-economy hubs, and these areas became far more educated. Hence previously low-educated places did not catch up with previously high-educated places like Oxford. All areas became far more highly-educated, such that the geographic distribution spread out and shifted to the right. In the mid-1990s, there were eight local authorities with more than one-third degree holders; by 2020, that number had increased to 298 (or 81% of all local authorities). There was, and there continues to be, a clustering of highly-educated people in certain places, but there are now far more of them than in the past. In 1996, the only borough that had more than 50% of its adult population degree-educated was the City of London. By 2020, in a quarter of all local authorities (91), a majority of adult residents held degrees. Finally, the distribution of house prices has evolved quite similar to education, but with more positive outliers. Houses have become considerably more expensive on average and more geographically varied; a small number of places—mostly in London—are now far more expensive than the average.

Far greater variation exists today in terms of education levels and age profiles; age and edu-

Figure 7: Socio-Economic Characteristics of Local Authority Districts, 1994-6 and 2020

Notes: The 1994-6 local authorities are mapped onto 2020 boundaries. Household income and house prices are inflation adjusted.

cation levels have become more important in describing places. The number of local authorities with more than one-third degree holders *and* more than 10% young people has grown eight-fold from six local authorities in the mid-1990s to 47 in 2020. This has led to a large rise in the number of areas with a majority of "universalists" (Bornschier et al. 2021)—people who are both highly socially liberal and highly supportive of redistribution—and a large number of areas where this group is in a minority. Combined with the concentration of support for redistribution at the individual level among highly-educated and young people, this means that age and education have also become much more important in describing how much a place supports redistribution and welfare. We document this in the following section.

The New Geographic Preference Divide

Figure 8 plots our MRP-estimated public opinions by local authority across the two time periods against the local authority socio-economic covariates from Figure 7, with the addition of population density. From 1994-6 the relationship between local authority household income and welfare

Figure 8: Social Policy Support and Socio-Economic Characteristics Across Local Authority Districts

Notes: Predicted local authority estimates based on 2020 YouGov data and 1994-6 BSAS data. Each dot represents a local authority. The colored lines show loess regressions. The dashed horizontal lines show median preferences (combined measure). Public opinion measures are normalized to the 0-1 interval. Higher preference values indicate more support for welfare and redistribution. Household income and house prices are inflation-adjusted.

opinions was downward-sloping over most of the data. Only a few high-income local authority outliers had high levels of support. This pattern reflects the traditional geographic coalition for redistributive policies: low-income areas were, in general, most supportive. By 2020, however, the relationship was mostly flat. Across most areas there was no relationship between average incomes and support, although again a few high-income areas stood out. The changing pattern for unemployment is even starker. In the mid-1990s, areas with the highest level of unemployment were much more supportive of welfare. Three decades later there was little or no relationship with local authority unemployment levels—with the slight exception of the redistribution scale.

By contrast, the education and age composition of local authority districts was far more predictive of social policy preference in 2020. In 1994-6 only a handful of places were highly educated, and within that small subset the relationship was upward-sloping. Otherwise, education was not very predictive of support. By 2020 an area's education had become an extremely powerful indicator of support. Meanwhile, there was a weakly positive relationship between the proportion of an area's population aged 18-24 and support in 1994-6, which became far stronger in the latter period. In terms of population density the relationships became slightly more positive over time too, but remained mostly flat; only the most urbanized areas stood out as more supportive. Finally, the pattern for house prices is quite similar to that for incomes. In the 1990s, places with lower house prices were a bit more supportive, especially of redistribution. But in 2020, there was little relationship, other than a few very expensive places being most supportive.

Implications for Welfare State Coalitions

These patterns have several implications for electoral coalitions in favor of social policies. The characteristics of local authorities that are necessary to win a political majority in favor of welfare have changed. For left-wing parties, we can think of this as winning all areas up to the median. In Figure 8, this means winning all areas from the top of each panel down to the dashed horizontal line, which indicates the median local authority in terms of social policy support. In the mid-1990s, "core" areas with the highest support had low incomes, high unemployment, and lower house prices. A party could secure a majority by winning all of these core areas together with places that had moderate incomes and unemployment rates, which were mostly low-educated and young to moderately young. These are the geographic contours of the traditional social democratic coali-

23

Figure 9: Social Policy Support and Leave Share in Brexit Referendum Across Local Authority Districts

tion. By 2020, however, the landscape of preferences had changed. The "core" areas of high support are now highly heterogeneous in terms of income, unemployment, and house prices but are overwhelmingly likely to be highly-educated and young. Winning a majority by expanding this to median-support areas means capturing a substantial number of areas with low to moderate education, in addition to places with high education; the same is true of age.

This also means that the geographic divide over welfare in 2020 maps closely onto the Brexit divide, an issue that divided socially liberal and conservative citizens. Figure 9 plots a local authority's leave vote share in the 2016 referendum against our public opinion estimates. In the 1994-6 period, areas that would become least likely to support Brexit showed the highest support for redistribution and welfare. But there was little relationship among areas with stronger support for leaving the EU. A substantial fraction of places that later went on to support leaving the EU were as supportive of redistribution as the least leave-supporting places. By 2020, the relationship was stronger across the board: local authorities with the highest leave vote share were the least supportive of redistribution and welfare, and none of those places showed support that was anywhere close to local authorities with the highest remain vote share. This is, of course, strongly linked to age and education, since support for remaining in the EU was strongest in the youngest and most educated places such as university towns. Geographically, this also suggests that aver-

Figure 10: Changes in Support for Welfare and Educational Attainment by Local Authority District, 1994-6 to 2020

Notes: Change in support from 1994-6 to 2020. Each dot represents a local authority. The colored lines show loess regressions. The 1994-6 data uses 2020 boundaries. Public opinion measures are normalized to the 0-1 interval.

age opinions on redistribution and welfare are now strongly correlated with social conservatism and opinions on "second-dimension" issues such as immigration and Euroscepticism.

Another way to think about these patterns is in terms of the debate about so-called "left-behind" areas. Recently the Conservative party has made inroads into traditional Labour-supporting areas like rural North-East England, with lower educational attainment and high leave support in the Brexit vote. Some analysts have suggested that this will place pressure on the Tories to embrace redistribution, or what is sometimes termed *levelling up*: "an expression of a realignment in British politics with the Conservatives presenting themselves as the new party of redistribution" (Jennings, McKay and Stoker 2021, p. 302). But in fact, it is precisely in the most "left-behind" places that support for redistribution and welfare has either fallen or risen the least. Figure 10 plots the change in support for welfare and redistribution against changes in the share of adults with degrees from 1994-6 to 2020, by local authority. The relationship is weakly positive across all three preference dimensions. In "left-behind" areas with little rise in higher education attainment, support for welfare and redistribution either fell or rose only slightly. But in places that experienced large educational upgrading, support grew much more. Because support for redistribution is now so strongly associated with the educational divide—and, by extension, an area's social conservatism—"left-behind" areas have polarized from more successful areas, becoming

less rather than more supportive. This in one explanation why the Conservative party focused on infrastructure, education, and job creation instead of traditional transfer policies in their levelling-up program as an electoral strategy to capture (and keep) votes in left-behind areas.

Electoral Challenges for Parties of the Left

How does this affect the ability of political parties to form a pro-redistribution geographic coalition? We now turn to the relationship between party vote shares and opinions on redistribution across local authorities in both periods. British local government elections take place at least every four years. Since not all local government elections are held at the same time, we combine election results from several years to match the two time periods for which we have public opinion data. For the 1994-6 period, we use election data from local elections held in 369 local authorities in England, Wales, and Scotland in 1994, 1995, and 1996. For 2020, we use results from local elections held in 293 local authorities in England in 2018 and 2019 (94% of all English local authorities).[12]

Figure 11 plots the relationship between average party vote shares and our estimated preferences for welfare and redistribution by local authority in both periods. Despite the geographic realignment of support for redistribution and welfare, for the most part there has *not* been a similar realignment of party support. The Conservative and Labour parties are still representing anti- and pro-redistribution areas, respectively. In both periods across most of the range of the data, the relationships were negative and positive, respectively. The deservingness dimension of the welfare state today is more predictive of an area's support for Labour than it was in the mid-1990s. Patterns for the Green party are similar to Labour, while the centrist Liberal Democrats has shifted over time from a largely anti-redistribution party to a mildly pro-redistribution one. This again shows that discussion of realignment in the UK, with the Conservative party allegedly representing "left-behind" areas that demand more redistribution, is premature. In fact, the party represents a coalition of areas with similar levels of support as in the 1990s, and there is no evidence that the party's new geographic base implies greater pressure for them to redistribute.

As we showed in Figure 5, the overall geographic distributions of social policy support changed very little between 1994-6 and 2020. But compared to the 1990s, today Labour is winning much

---

[12]We do not have data on local elections in Scotland and Wales as they were held in May 2017.

Figure 11: Social Policy Support and Local Election Vote Shares Across Local Authority Districts

Notes: Dots represent local authorities. The colored lines show loess regressions. The 1994-6 period uses election results from 369 local authorities in England, Wales, and Scotland, mapped onto 2020 local authority boundaries. The 2020 period uses 2018 and 2019 election results for 293 local authorities in England. Public opinion measures are normalized to the 0-1 interval.

lower vote shares across those distributions, as indicated by the left-ward shift in the mass of points in Figure 11. The opposite is true of the Conservative Party. Figure 12 shows part of the explanation for this by plotting Labour and Conservative party vote shares against district socio-demographic characteristics. Recall that in the mid 1990s, a geographic majority for redistributive policies consisted of mostly low-educated districts with low to medium incomes and unemployment rates and with a wide age distribution. In that period, Labour was winning in precisely those areas. It did much better in low-income, high-unemployment and low- to moderate-education areas, and out-performed the Conservative Party across most of the age spectrum. But in 2020 things had changed. Now a geographic coalition consisted of places across a much wider spectrum of incomes and unemployment rates, high to moderate education levels, and low to moderate numbers of young people. Although Labour still did better in the lowest-income areas in 2020, its performance in low-educated areas collapsed; it also did poorly in moderately-educated places. By contrast, the Conservatives did as well or better than Labour in all except the most highly-

Figure 12: Party Vote Shares and Socio-Economic Characteristics Across Local Authority Districts

Notes: Dots represent local authorities. The colored lines show loess regressions. The 1994-6 period uses election results from 369 local authorities in England, Wales, and Scotland, mapped onto 2020 local authority boundaries. The 2020 period uses 2018 and 2019 election results for 293 local authorities in England. Household income is inflation adjusted.

educated areas. They also outperformed Labour by a very wide margin in seats with low numbers of young people and did at least as well in average areas; Labour won only the youngest places.

In geographic terms, the realignment of political competition over welfare along "second-dimension" lines, and the associated decisiveness of moderately-educated and moderately-young areas, makes it very difficult for Labour to win a majority in the legislature for redistributive policies, even with almost the same aggregate support for these policies as in the 1990s. Areas close to the median support are often opposed to the EU and immigration and are more socially conservative. Even if they are closely aligned with Labour's views on redistribution, they are much further apart on these non-economic issues, which now form a core part of political campaigns. Relatively few places have high numbers of graduates or high numbers of young people—which

Labour wins by wide margins—but it is losing moderately-educated and moderately-young areas. The rise of "second-dimension" politics is not itself the core problem since popular support for welfare and redistribution has not changed over time. Instead, the same level of popular support does not translate as easily into a geographic majority because votes in favor of welfare and redistribution are, from the perspective of the left, inefficiently distributed across areas.

Conclusion

We have argued that the spatial clustering of voters by age and education shapes redistributive politics. When elections are held under majoritarian rule, this undermines the ability of left-wing parties to form redistributive coalitions. We provided evidence for our argument by measuring support for redistribution and deservingness across British local authorities, using multilevel regression and post-stratification applied to historical survey data from the mid-1990s and an original representative survey from 2020. We showed that an area's average income and unemployment rate were highly predictive of social policy support in the mid-1990s; today, this is no longer the case. Instead, support for the welfare state has become concentrated in areas dominated by young and highly educated voters—mostly urban centers and university towns. Such socially-liberal "universalists" have grown in number and are now the demographic core of support for welfare in the UK. Contemporary social policy preferences closely resemble patterns associated with so-called "second-dimension" issues like immigration and Brexit. The rise of political competition along this dimension, together with the spatial concentration of univeralists, makes assembling a geographic coalition for redistribution and welfare more difficult than in the 1990s, even though aggregate support is virtually identical. Despite a similar number and *majority of voters* supporting such policies, it has become harder to secure a *majority of places* in favor of them.

We then linked these geographical dynamics to voting patterns and showed that, despite considerable change in the type of areas that support redistribution, the Labour and Conservative parties continue to be pro- and anti-redistribution parties, respectively, in terms of the areas they represent. At the same time, the arrival of new "red wall" areas—that is, constituencies that historically tended to support the Labour party—in the Conservative coalition has *not* led to the party

representing places that demand greater redistribution, because those places are relatively old and low-educated, meaning that they no longer form the geographical core of support for social policy.

Our findings have several implications for redistributive politics, the future of left-wing parties, and place-based socio-economic policies. First, our paper provides a new explanation for why majoritarian electoral systems make redistribution more difficult to achieve. While majoritarian democracies face increasing difficulties assembling geographic coalitions for redistribution due to the geographic concentration of universalists, countries with proportional representation (PR) rules are less likely to be affected by these dynamics because political competition is less dependent on the spatial distribution of socio-economic groups. Political conflict in majoritarian democracies tends to unfold along geographic lines, thus undermining traditional redistributive policies; under PR, conflict tends to be structured along class lines, pitting highly-educated and high-income groups against low-educated and low-income groups regardless of place. But if economic opportunities are more likely to cluster in urban compared to rural areas, PR electoral rules might amplify the voices of the urban coalition and implement pro-urban policies while neglecting rural interests. This could lead to further electoral backlash against urban elites.

Second, the patterns we uncover in this paper may lead to a restructuring of political competition in majoritarian democracies such as the Canada, France, the UK, and the United States. In these countries, politics may become dominated by an urban coalition of high-skilled workers and low-income voters demanding conventional redistribution, opposing another coalition of rural and suburban voters preferring intra-regional redistribution away from knowledge-economy centers (Ansell and Gingrich 2021; Beramendi 2012). Our findings provide clear evidence that such a realignment is plausible—or even under way already. This complicates the task of reducing regional inequalities through transfer-based redistributive policies because "left-behind" areas, many of which are also socially conservative, are not strongly supportive of government intervention through redistribution and welfare provision. Any politically sustainable project to improve such areas may require redistribution through place-based policies—for example spending on local infrastructure or job creation as prominently discussed in British politics under the heading of "levelling up"—rather than via transfer-based welfare programs. The Tories have already adopted such policies as an electoral strategy, hoping to secure votes in former Labour party strongholds.

Third, our findings suggest that researchers may need to consider preferences for redistribu-

30

tion and welfare separately from preferences over other forms of economic policy, which may be less tied to age and education. The same high-income, socially liberal voters that support welfare policies may not be as supportive of tax-and-spend issues that fit more conventionally onto the economic "first dimension." This is not captured by our survey questions but is reflected in the changing nature of the welfare state over recent decades, with rising for social investment policies by well-educated middle class voters (Gingrich and Häusermann 2015). Nonetheless, our findings—alongside recent work by Bornschier et al. (2021)—also call into question whether politics in some European democracies is becoming one-dimensional, pitting the universalist, socially liberal middle classes against particularist, socially conservative voters whose views on redistribution and social issues increasingly align. Our paper contributes to debates about the political viability of redistribution and welfare in this potentially more unidimensional environment.

References

Abou-Chadi, Tarik and Simon Hix. 2021. "Brahmin Left versus Merchant Right? Education, class, multiparty competition, and redistribution in Western Europe." *British Journal of Sociology* 72(1):79–92.

Abrams, Samuel J. and Morris Fiorina. 2012. ""The Big Sort" That Wasn't: A Skeptical Reexamination." *PS: Political Sciences and Politics* 45(2):203–210.

Adler, David and Ben Ansell. 2020. "Housing and Populism." *West European Politics* 43(2):344–365.

Ansell, Ben and Jane Gingrich. 2021. Concentration and Commodification: The Political Economy of Postindustrialism in America and Beyond. In *The American Political Economy: Politics, Markets, and Power*, ed. Jacob S. Hacker, Alexander Hertel-Fernandez, Paul Pierson and Kathleen Thelen. Cambridge: Cambridge University Press pp. 375–406.

Attewell, David. 2021. "Redistribution attitudes and vote choice across the educational divide." *European Journal of Political Research*, forthcoming.

Baccini, Leonardo and Stephen Weymouth. 2021. "Gone For Good: Deindustrialization, White Voter Backlash, and US Presidential Voting." *American Political Science Review* 115(2):550–567.

Bell, Torsten and Laura Gardiner. 2019. "My Generation, Baby: The Politics of Age in Brexit Britain." *The Political Quarterly* 90(S2):128–141.

Beramendi, Pablo. 2012. *The Political Geography of Inequality: Regions and Redistribution*. Cambridge Studies in Comparative Politics Cambridge: Cambridge University Press.

Boix, Carles. 2019. *Democratic Capitalism at the Crossroads: Technological Change and the Future of Politics*. Princeton: Princeton University Press.

Bornschier, Simon, Silja Häusermann, Delia Zollinger and Céline Colombo. 2021. "How "Us" and "Them" Relates to Voting Behavior—Social Structure, Social Identities, and Electoral Choice." *Comparative Political Studies,* forthcoming.

BSAS. 2021. *New Values, New Divides?* NatCen Social Research: British Social Attitude Survey.

Bullock, John. 2021. "Education and Attitudes toward Redistribution in the United States." *British Journal of Political Science* 51(3):1230–1250.

Bürkner, Paul-Christian. 2018. "Advanced Bayesian Multilevel Modeling with the R Package brms." *The R Journal* 10(1):395–411.

Caughey, Devin, Tom O'Grady and Christopher Warshaw. 2019. "Policy Ideology in European Mass Publics, 1981–2016." *American Political Science Review* pp. 1–20.

Cavaillé, Charlotte and Kris-Stella Trump. 2015. "The Two Facets of Social Policy Preferences." *The Journal of Politics* 77(1):146–160.

Cho, Wendy Tam, James Gimpel and Iris Hui. 2013. "Voter Migration and the Geographic Sorting of the American Electorate." *Annals of the Association of American Geographers* 103(4):856–870.

Cramer, Katherine J. 2016. *The Politics of Resentment: Rural Consciousness in Wisconsin and the Rise of Scott Walker*. Chicago: University of Chicago Press.

Döring, Holger and Philip Manow. 2017. "Is Proportional Representation More Favourable to the Left? Electoral Rules and Their Impact on Elections, Parliaments and the Formation of Cabinets." *British Journal of Political Science* 47(1):149–164.

Enos, Ryan D. 2017. *The Space between Us: Social Geography and Politics*. Cambridge: Cambridge University Press.

Gallego, Aina, Franz Buscha, Patrick Sturgis and Daniel Oberski. 2016. "Places and Preferences: A Longitudinal Analysis of Self-Selection and Contextual Effects." *British Journal of Political Science* 46(3):529–550.

Gelepithis, Margarita and Marco Giani. 2020. "Inclusion without Solidarity: Education, Economic Security, and Attitudes toward Redistribution." *Political Studies,* forthcoming.

Gidron, Noam and Peter Hall. 2017. "The politics of social status: economic and cultural roots of the populist right." *British Journal of Sociology* 68(S1):S57–S84.

Gimpel, James G., Nathan Lovin, Bryant Moy and Andrew Reeves. 2020. "The Urban–Rural Gulf in American Political Behavior." *Political Behavior* 42(4):1343–1368.

Gingrich, Jane and Silja Häusermann. 2015. "The Decline of the Working-Class Vote, the Reconfiguration of the Welfare Support Coalition and Consequences for the Welfare State." *Journal of European Social Policy* 25(1):50–75.

Green, Jane and Roosmarijn de Geus. 2021. "Red Wall, Red Herring? Economic Insecurity and Voting Intention in Britain." *Nuffield Politics Research Centre Report.* .

Hanretty, Chris. 2020. "An Introduction to Multilevel Regression and PostStratification for Estimating Constituency Opinion." *Political Studies Review* 18(4):630–645.

Hanretty, Chris, Benjamin E. Lauderdale and Nick Vivyan. 2018. "Comparing Strategies for Estimating Constituency Opinion from National Survey Samples." *Political Science Research and Methods* 6(3):571–591.

Häusermann, Silja and Hanspeter Kriesi. 2015. What Do Voters Want? Dimensions and Configurations in Individual-Level Preferences and Party Choice. In *The Politics of Advanced Capitalism*, ed. Hanspeter Kriesi, Herbert Kitschelt, Pablo Beramendi and Silja Häusermann. Cambridge: Cambridge University Press pp. 202–230.

Häusermann, Silja, Thomas Kurer and Hanna Schwander. 2015. "High-Skilled Outsiders? Labor Market Vulnerability, Education and Welfare State Preferences." *Socio-Economic Review* 13(2):235–258.

Hjorth, Frederik. 2016. "Who benefits? Welfare chauvinism and national stereotypes." *European Union Politics* 17(1):3–24.

Inglehart, Ronald. 1990. *Culture Shift in Advanced Industrial Society*. Princeton, N.J.: Princeton University Press.

Iversen, Torben and David W. Soskice. 2019. *Democracy and Prosperity: Reinventing Capitalism Through a Turbulent Century*. Princeton: Princeton University Press.

Jennings, Will, Lawrence McKay and Gerry Stoker. 2021. "The Politics of Levelling Up." *The Political Quarterly* 92(2):302–310.

Jusko, Karen Long. 2017. *Who Speaks for the Poor?: Electoral Geography, Party Entry, and Representation*. Cambridge Studies in Comparative Politics Cambridge: Cambridge University Press.

Kastellec, Jonathan P., Jeffrey R. Lax and Justin H. Phillips. 2010. "Estimating state public opinion with multi-level regression and poststratification using R." *Unpublished Manuscript, Princeton University* .

Kitschelt, Herbert P. and Philipp Rehm. 2022. "Polarity Reversal: The Socioeconomic Reconfiguration of Partisan Support in Knowledge Societies." *Politics & Society* pp. 1–47.

Kurer, Thomas. 2020. "The Declining Middle: Occupational Change, Social Status, and the Populist Right." *Comparative Political Studies* 53(1):1798–1835.

Kuziemko, Ilyana, Ryan W. Buell, Taly Reich and Michael I. Norton. 2014. ""Last-Place Aversion": Evidence and Redistributive Implications." *The Quarterly Journal of Economics* 129(1):105–149.

Marshall, John. 2016. "Education and voting Conservative: Evidence from a major schooling reform in Great Britain." *Journal of Politics* 78(2):382–395.

Marshall, John. 2019. "The anti-Democrat Diploma: How High School Education Decreases Support for the Democratic Party." *American Journal of Political Science* 61(1):67–83.

Maxwell, Rahsaan. 2019. "Cosmopolitan Immigration Attitudes in Large European Cities: Contextual or Compositional Effects?" *American Political Science Review* 113(2):456–474.

McArthur, Daniel. 2022. "Why are the highly educated more sympathetic towards welfare recipients?" *European Journal of Political Research* forthcoming.

McKee, Seth C. 2008. "Rural Voters and the Polarization of American Presidential Elections." *PS: Political Science & Politics* 41(1):101–108.

Mendelberg, Tali, Katherine T. McCabe and Adam Thal. 2017. "College Socialization and the Economic Views of Affluent Americans." *American Journal of Political Science* 61(3):606–623.

Mewes, Jan and Steffen Mau. 2013. "Globalization, socio-economic status and welfare chauvinism: European perspectives on attitudes toward the exclusion of immigrants." *International Journal of Comparative Sociology* 54(3):228–245.

O'Grady, Tom. 2022. *The Transformation of British Welfare Policy: Politics, Discourse and Public Opinion*. Oxford: Oxford University Press.

Pearce, Nick and Joe Chrisp. 2021. "The Age Divide in UK Politics." *Institute for Policy Research Working Paper, University of Bath.* .

Reeskens, Tim and Wim van Oorschot. 2012. "Disentangling the 'New Liberal Dilemma': On the relation between general welfare redistribution preferences and welfare chauvinism." *International Journal of Comparative Sociology* 53(2):120–139.

Rodden, Jonathan. 2010. "The Geographic Distribution of Political Preferences." *Annual Review of Political Science* 13(1):321–340.

Rodden, Jonathan. 2019. *Why Cities Lose: The Deep Roots of the Urban-Rural Political Divide*. New York: Basic Books.

Rodríguez-Pose, Andrés. 2018. "The Revenge of the Places That Don't Matter (and What to Do About It)." *Cambridge Journal of Regions, Economy and Society* 11(1):189–209.

Rueda, David. 2018. "Food Comes First, Then Morals: Redistribution Preferences, Parochial Altruism, and Immigration in Western Europe." *Journal of Politics* 80(1):225–239.

Sands, Melissa. 2017. "Exposure to inequality affects support for redistribution." *Proceedings of the National Academy of Sciences* 114(4):663–668.

Warshaw, Christopher and Jonathan Rodden. 2012. "How Should We Measure District-Level Public Opinion on Individual Issues?" *Journal of Politics* 74(1):203–219.

Wiedemann, Andreas. 2023. "Redistributive Politics Under Spatial Inequality." *Journal of Politics* .

Biographical Statement

Tom O'Grady is an Associate Professor of Political Science at University College London, UK. Andreas Wiedemann is an Assistant Professor of Politics and International Affairs at Princeton University and the Princeton School of Public and International Affairs, NJ, USA.

Supplementary Appendix

A Comparison of YouGov Survey to 2020 BSAS

Figure A.1: Mean Support for Welfare and Redistribution (3 Scales), 2020 YouGov compared to 2020 BSAS

Notes: Shows Mean support and 95% confidence intervals. All scales range from 1-5, where higher values indicate more left-wing opinions

Figure A.2: Mean Support for Welfare and Redistribution (5 Survey Questions), 2020 YouGov compared to 2020 BSAS

Notes: Mean support and 95% confidence intervals. All scales range from 1-5, where higher values indicate more left-wing opinions

For the three scales used in the paper, average support was almost identical in the two surveys (Figure A.1). The scales range from 1-5; all differences are therefore substantively very small.

Looking instead at the five individual survey items that make up these scales (), we find a meaningful difference between the two surveys only for *redistrb*. YouGov respondents in 2020 were more supportive of redistribution than those of the BSAS in 2019 by 0.23 points on the 5-point scale, about one-fifth of a standard deviation. For the other four items, differences were negligible.

Subgroup differences were also almost identical between the two surveys (Figure A.3). For both the combined and deservingnes dimensions, as well as for redistribution and education, the age and education differences are virtually indistinguishable. Only for redistribution and age do we find some difference: the gap between young and old is a little bit larger in the YouGov survey than in the BSAS.

Figure A.3: Mean Support for Welfare by Subgroup, 2020 YouGov compared to 2020 BSAS

Notes: Shows mean support and 95% confidence intervals. All scales range from 1-5, where higher values indicate more left-wing opinions

B Data Sources

Table B.1: Variable Description and Data Sources

Variable	Years	Data Sources
Socio-economic variables, local authority level		
Population age-sex groups	1991, 2001; 2019	ONS: Nomis Annual Population Survey
Benefit claimant counts by sex and age	1996; 2020	ONS: Nomis claimant counts
Share population aged 18-24 (%)	1994-6; 2019	ONS: Nomis population estimates
Share population aged 65 and older	1994-6; 2019	ONS: Nomis population estimates
Total gross disposable household income	1997; 2017	ONS: Regional accounts
Average amount of benefits received, per person	1997; 2017	ONS: Regional accounts
Unemployment rate	1996; 2020	ONS: M01 Regional labour market
Population density (people per square km)	1994-6; 2019	ONS: Nomis population estimates
Average house sales prices	1994-6; 2020	HM Land Registry: UK House Price Index; Registers of Scotland: House price statistics
Share adults educated to degree level or higher	1991, 2001, 2019	ONS: 1991 Census data; ONS: Nomis 2001 Census Data; ONS: Nomis Annual Population Survey; Scottish Registrar General's Report on 2001 Census
Socio-economic variable, regional authority level		
Annual change in gross value added	1996; 2020	ESCOE. Regional Nowcasting
Election variables, local authority level		
Local election party vote shares for England, Scotland, Wales	1994-6	The Election Center
Local election party vote shares for England	2018; 2019	The Election Center
Brexit referendum vote shares	2016	Electoral Commission

Notes: All variables are for local authorities in England, Wales, and Scotland unless otherwise noted. Estimates of UK regional quarterly GVA based on a model by Koop, Gary, Stuart McIntyre, James Mitchell and Aubrey Poon. 2020. "Regional output growth in the United Kingdom: More timely and higher frequency estimates from 1970." *Journal of Applied Econometrics* 35(2):176–197.

Construction of House Price Measure

House price data at the local authority level are not consistently available for the UK. While the ONS "UK house price data" exists for England since 1970 and Wales since 1995, it exists for Scotland only since 2004. The Registers of Scotland, however, have produced their own house price dataset starting in 1993, but their data series does not match perfectly with the Scottish series from the ONS. Our solution is to use 2004 as a benchmark year—where the Scottish time series in the ONS UK house price data starts—and adjust the longer time series from the Registers of Scotland by house price value in a given local authority in 2004. For 2004, the longer time series from the Registers of Scotland will match the shorter time series from the ONS UK house price data. Using that inflation factor, we deflate the pre-2004 data from the Registers of Scotland and take the average of the 1994-6 period based on the adjusted time series. For 2020 we use data from the ONS UK house price index.

Years of Data Used in Our Models

When assembling data our models, the uneven availability of data required us to be pragmatic about which dates to use these socio-economic series for in our models; we used the nearest available date in all cases. The closest-matching dates for our data on the benefit claimant count, the share of population aged 18-24, the share of population aged 65 and older, gross value added, and population density are 1996 and 2019 respectively. For mean gross disposable household income and mean amount of benefits received per person, the closest-available dates are 1997 and 2017 respectively (no data are available prior to 1997 in these cases). It is reasonable to assume that, for example, mean gross disposable income in a local authority in 2017 is a very close proxy for income in 2020. For the proportion of a local authority's adult population educated to degree level, we use 2019 data in the latter period. Prior to the 2000s this data is available only in decennial census years. We therefore use for the 1994-6 period an average of the 1991 and 2001 census data. Given the near-linear increase in degree-level education over the period, this average should again be a good proxy for the 1994-6 period.

C Comparison of Bayesian and Non-Bayesian Estimates

Figures C.1 and C.2 compare our main estimates, produced with a Bayesian approach, to estimates from a non-Bayesian multilevel model estimated with R's `lmer` function. Both show estimates for our combined scale. Each dot represents a local authority area. The figures show that the results from our Bayesian approach are virtually indistinguishable from a standard non-Bayesian multilevel model.

Figure C.1: Comparison of Bayesian and non-Bayesian Estimates, 1996 BSAS data (Combined scale)

Figure C.2: Comparison of Bayesian and non-Bayesian Estimates, 2020 YouGov data (Combined scale)

D Uncertainty of Bayesian MRP Estimates

Figure D.1: MRP Estimates by Local Authority Districts with 95% Credible Intervals

Notes: The black line shows posterior means with 95% credible intervals displayed in gray.
Higher values indicate more support for social policies.

Figure D.2: Average MRP Estimates by Local Authority Districts with 95% Credible Intervals

Notes: This figure shows the average mean and 95% credible intervals across all local authority districts for each social policy dimension in 1994-6 and 2020.

E   Maps of All Preference Dimensions

Figure E.1: Predicted Social Policy Support by Local Authority Districts, Great Britain

(a) 1994-6 Public Opinion

(b) 2020 Public Opinion

Notes: Predicted local authority estimates based on 2020 survey data and 1994-6 BSAS data. Social policy estimates are normalized to the 0-1 interval. Higher values indicate more support for social policies. Boundaries differed in 1994-6 and 2020. We have no estimates for the City of London due to lack of covariate data.