

The Polarizing Effect of Partisan Echo Chambers

SARA B. HOBOLT *London School of Economics and Political Science, United Kingdom*

KATHARINA LAWALL *University of London, United Kingdom*

JAMES TILLEY *University of Oxford, United Kingdom*

We are witnessing increasing partisan polarization across the world. It is often argued that partisan “echo chambers” are one of the drivers of both policy and affective polarization. In this article, we develop and test the argument that the political homogeneity of people’s social environment shapes polarization. Using an innovative, large-scale pre-registered “lab-in-the-field” experiment in the United Kingdom, we examine how polarization is influenced by partisan group homogeneity. We recruit nationally representative partisans and assign them to discuss a salient policy issue, either with like-minded partisans (an echo chamber) or in a mixed-partisan group. This allows us to examine how group composition affects polarization. In line with our expectations, we find that partisan echo chambers increase both policy and affective polarization compared to mixed discussion groups. This has important implications for our understanding of the drivers of polarization and for how out-group animosity might be ameliorated in the mass public.


INTRODUCTION


Political conflict is central to democratic societies. However, democracy suffers when such conflict solidifies and political identities crystallize into polarized groups who are unwilling to engage respectfully with each other. In recent decades, much attention has been paid to the increasing partisan polarization in American politics (Layman, Carsey, and Horowitz 2006; Mason 2015; 2018). Of particular concern is that partisan polarization is not just, or even primarily, about policy disagreement, but manifests itself as affective polarization: an emotional attachment to the in-group party and hostility toward the out-group party (Green, Palmquist, and Schickler 2004; Iyengar et al. 2019; Iyengar, Sood, and Lelkes 2012; Iyengar and Westwood 2015). Recent work has shown that affective polarization is not unique to the United States, but is also found in other Western democracies (Gidron, Adams, and Horne 2020; Hartevelde 2021; Hobolt, Leeper, and Tilley 2021; Reiljan 2020; Wagner 2021). The negative consequences of such polarization include out-group prejudice and discrimination (Huber and Malhotra 2017; Iyengar et al. 2019), a lack of


willingness to engage in democratic deliberation, and even reduced trust in democratic institutions and support for basic democratic principles (Hetherington and Rudolph 2015; Kingzette et al. 2021; Layman, Carsey, and Horowitz 2006).¹

What explains affective polarization? One prominent explanation is the existence of partisan “echo chamber” environments in which people encounter only the opinions and beliefs of like-minded partisans. Some highlight how social media may contribute to such political echo chambers (Sunstein 2018; Terren and Borge-Bravo 2021; Toernberg et al. 2021). Equally, this process could be related to partisan sorting of the type we have seen in the United States, along regional, racial, and religious lines, that reinforces partisan identities (Mason 2015; 2018). Either way, by surrounding ourselves, in real life and online, with others who share our perspectives and opinions about the world, we amplify tribalism and exacerbate polarization.

Yet while much has been written on the role of echo chambers as a cause of affective polarization, few studies actually test the causal effect of group homogeneity on polarization in a social context. In this article, we argue that since affective polarization is rooted in social attachments to political in-groups and hostility toward out-groups, the drivers of such polarization can best be explored in a social setting. Specifically, we build on the literature on social identities and intergroup contact to argue that someone’s closeness to their in-party, as well as their animosity toward the out-party and partisans, are influenced by the homogeneity of their political discussion partners within their real-life social environment (Allport 1954; Levendusky, Druckman, and

Corresponding author: Sara B. Hobolt , Sutherland Chair in European Institutions, Department of Government, London School of Economics and Political Science, United Kingdom, s.b.hobolt@lse.ac.uk.

Katharina Lawall , Postdoctoral Fellow, Centre for the Politics of Feelings, University of London, United Kingdom, katharina.lawall@sas.ac.uk.

James Tilley , Professor of Politics, Department of Politics and International Relations, University of Oxford, United Kingdom, james.tilley@politics.ox.ac.uk.

Received: September 07, 2022; revised: April 25, 2023; accepted: October 19, 2023.

¹ Although see Broockman, Kalla, and Westwood (2023), who find no effect of affective polarization on democratic norms.

McLain 2016; Pettigrew and Tropp 2006; Visser and Mirabile 2004). Greater partisan homogeneity is likely to increase polarization, not just in terms of people's policy stances, but also their out-group animosity.

To test this argument, we conducted an innovative online “lab-in-the-field” experiment in the United Kingdom. We recruited partisans of the two major British parties, Labour and the Conservatives, to take part in group discussions about a divisive political issue (a new government policy on immigration and asylum). Participants were randomly assigned to discuss the issue with either fellow partisans or a mixed cross-party group. By surveying participants before and after these group discussions, our design allows us to examine the causal effect of group composition on both policy and affective polarization. In line with previous research, we find that policy polarization is higher in homogeneous discussion groups compared to cross-party discussion groups (Klar 2014; Mutz 2006). Importantly, we also find that echo chambers make people significantly more affectively polarized than do cross-partisan group discussions. These findings go further than other experimental studies in which discussion takes place exclusively in cross-partisan groups (Fishkin et al. 2021), or exclusively in echo chambers (Schkade, Sunstein, and Hastie 2010), or where the topic of discussion is one that encourages agreement (Levendusky and Stecula 2021; Santoro and Broockman 2022). Interestingly, we find that while the policy polarization effect occurs among both Conservative and Labour supporters, the impact on affective polarization is driven by Labour partisans. We suggest that supporters of an opposition party, in this case Labour, might be particularly susceptible to affective polarization when they find themselves in an echo chamber.

Our article makes three major contributions. First, we provide compelling evidence that concerns about the role of real-life echo chambers in a polarizing society are justified. Policy discussion in homogeneous groups, compared to mixed groups, produces not only greater policy disagreement, but also greater affective polarization. Second, these conclusions are based on an innovative lab-in-the-field experiment. While there is a large literature on affective polarization both in the United States and elsewhere, the vast majority of work relies on observational survey data. More recently, there has been a rise in important experimental work examining the effect of group composition on polarization (Groenlund, Herne, and Setälä 2015; Klar 2014; Levendusky and Stecula 2021; Santoro and Broockman 2022). Our experiment contributes to this growing literature. Importantly, it is designed to resemble a real political discussion, recruiting a representative sample of partisans to participate in live group discussions online about a divisive and salient issue, without imposing strong deliberative norms or emphasizing cross-partisan similarities (Baron et al. 2021; Esterling, Fung, and Lee 2021; Fishkin et al. 2021; Groenlund, Herne, and Setälä 2015; Levendusky, Druckman, and McLain 2016; Levendusky and Stecula 2021). Third, our findings have important wider implications for how to tackle affective polarization. We demonstrate that both

positive and negative partisan affect are exacerbated by discussion in homogeneous partisan groups. While the conclusion that echo chambers worsen affective polarization is concerning, our study also highlights the potential effectiveness of policy strategies that increase the partisan heterogeneity of discussion groups. Our findings, therefore, suggest that increasing cross-partisan discussion, even when debating controversial issues, can ameliorate affective polarization.

PARTISAN ECHO CHAMBERS AND AFFECTIVE POLARIZATION

Affective polarization is an increasing concern among scholars of U.S. politics. Numerous studies have shown that the U.S. public has become more divided along partisan and ideological lines in recent years (Hetherington 2009; Layman, Carsey, and Horowitz 2006; Mason 2018). This has been accompanied by affective polarization, most notably increasing animosity toward out-group partisans (Iyengar et al. 2019; Iyengar, Sood, and Lelkes 2012; Iyengar and Westwood 2015; Layman, Carsey, and Horowitz 2006; Mason 2015; 2018). While affective polarization is often rooted in policy disagreement, it is distinct from ideological polarization. The latter concerns the extremity of political views, whereas the former is focused on highly emotive, affective in-group favorability and out-group animosity across party lines (Iyengar, Sood, and Lelkes 2012; Mason 2015; 2018).

Recent comparative work has shown that Americans are not unique when it comes to strong positive in-party feelings and dislike of out-parties (Gidron, Adams, and Horne 2020; Hartevelde 2021; Kekkonen and Yl-Anttila 2021; Reiljan 2020; Torcal and Comellas 2022; Wagner 2021). Wagner (2021) convincingly demonstrates that the United States is unexceptional when it comes to affective polarization around the world, and Reiljan's (2020) comparative study concludes that partisans in the United States are less affectively polarized than their counterparts in Central Eastern and Southern Europe. The widespread nature of affective polarization is clearly of concern, since well-functioning democracies need citizens and politicians to engage respectfully with each other, even on controversial topics (Lipset 1959). High levels of mass affective polarization are likely to breed intolerance, political cynicism and reduced opportunities for collaboration and compromise (Hetherington and Rudolph 2015; Layman, Carsey, and Horowitz 2006; MacKuen et al. 2010). Some even argue that affective polarization can undermine support for democratic norms (Kingzette et al. 2021; McCoy, Rahman, and Somer 2018), although there is mixed evidence for this (Broockman, Kalla, and Westwood 2023; Santoro and Broockman 2022).

What has caused this rise in affective polarization? There are two common explanations given. On the one hand, it is argued that as partisan and ideological identities have become increasingly aligned—particularly with other group identities, such as race and religion—

the strength of partisan attachments has grown (Iyengar et al. 2019; Mason 2015; 2018). This decline of cross-cutting identities intensifies both emotional attachment to in-groups, but also hostility to the out-group as partisans find it easier to make generalized inferences about the “other side” (Mason 2015; 2018). On the other hand, it is argued that the elite political context is also important (Gidron, Adams, and Horne 2020; Iyengar et al. 2019). In particular, it has been shown that growing ideological polarization in the United States among political elites (Banda and Cluverius 2018; Rogowski and Sutherland 2016) and a polarized media environment have led to greater affective polarization in the mass public (Lelkes, Sood, and Iyengar 2017; Levendusky 2013; Suhay, Bello-Pardo, and Maurer 2018). Both of these approaches depend, to some extent, on the idea of people encountering increasingly homogeneous political views in partisan echo chambers. The alignment of social and political identities means that voters increasingly find themselves surrounded by “people like them.”

We can think of these echo chambers as producing both policy and affective polarization. Starting with the former, there is some experimental evidence that homogeneous interpersonal networks influence policy polarization. In particular, discussion among fellow partisans can lead individuals to agree more with their in-party policy position (Klar 2014; Schkade, Sunstein, and Hastie 2010; Visser and Mirabile 2004). There are a number of related explanations for this process. First, there is motivated reasoning: people overvalue information and arguments that support their own views and discount arguments from the other side (Nyhan and Reifler 2010). Second, there is social conformity: people adjust their policy opinions to fit their group norm (Festinger 1950; Sunstein 2009). People surrounded by others who share a particular attitude are socially rewarded for expressing views that reinforce that majority opinion (Schachter 1951). Third, attitudinal homogeneity within networks tends to increase an individual’s attitude strength as hearing similar views makes them less ambivalent about their own attitudes (Festinger 1950; Schkade, Sunstein, and Hastie 2010; Visser and Mirabile 2004). Conversely, discussion in mixed groups increases empathy through perspective taking, involving the active process of imagining oneself in another person’s shoes (Muradova 2021), which weakens policy polarization (Fishkin et al. 2021).

Yet the evidence on whether partisan group homogeneity leads to policy polarization is not clear-cut. Some studies conducted in deliberative settings, which encourage strong deliberative norms of respect and openness to other opinions, have found that, even in like-minded groups, discussion can lead to depolarization (Esterling, Fung, and Lee 2021; Groenlund, Herne, and Setala 2015). This work has greatly added to our understanding of how deliberation can reduce polarization, yet any depolarizing effect may also be caused by the way in which the deliberative setting seeks to promote cross-partisan understanding and balanced information processing. Indeed, it is argued that it is precisely the structuring of discussion in this deliberative way that

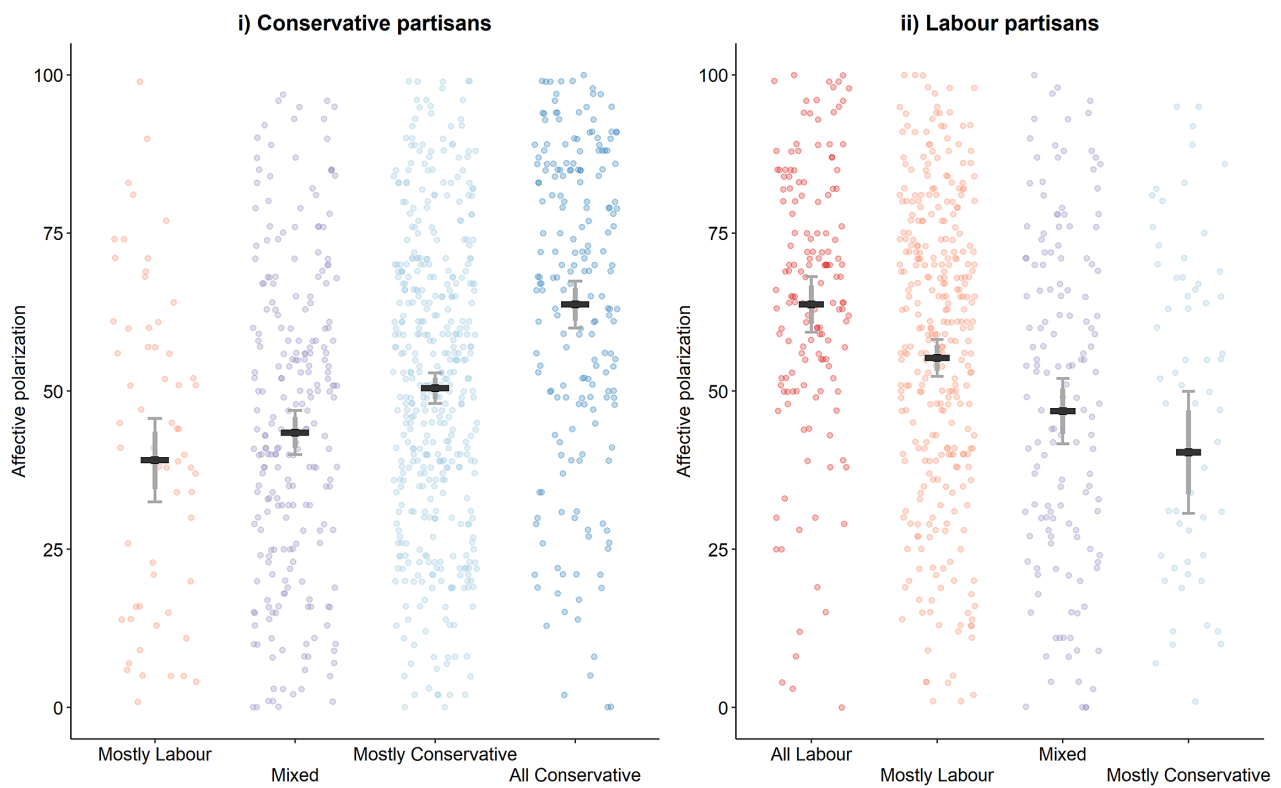
reduces opinion extremity in homogeneous groups (Strandberg, Himmelroos, and Groenlund 2019). Hence, in settings with less moderation and balance, which therefore more closely approximate real-life political discussion, we should find that people in homogeneous partisan groups are more polarized over policy. This leads to our first hypothesis about the role of the echo chamber on policy polarization.

H1: *Homogeneous discussion groups will increase policy polarization compared to heterogeneous discussion groups.*

Departing from most previous work, this article is primarily interested in the effect of partisan echo chambers on affective polarization, rather than policy polarization. We expect that similar mechanisms are at play, but we have only limited evidence about how the partisan homogeneity of networks affects affective polarization. Much of the interest in this question has focused on how social media fosters echo chambers, which may in turn cause affective polarization. It has been argued that social networking sites, such as Facebook and Twitter/X, encourage people to seek out networks of like-minded individuals. As a consequence, people are only exposed to information that reinforces their policy views (Gentzkow and Shapiro 2011; Pariser 2011; Sunstein 2018). Nonetheless, it is unclear how widespread echo chambers are on social media (Barbera 2020; Barbera et al. 2015; Nyhan et al. 2023) and, to the extent they do exist, there is no consensus as to the effects of them on in-group identities and out-group animosity. Some find that exposure to political views, and even partisan echo chambers, on social media does not have polarizing effects (Nyhan et al. 2023; Suhay, Bello-Pardo, and Maurer 2018), while others argue that social media usage increases perceptions of differences between partisan groups in ways that contribute to out-group prejudices and animosity (Settle 2018).

Beyond the study of social media effects, we have limited evidence on how discussion within the in-group or with the out-group has an impact on affective polarization. One important study that examines the effect of interpersonal group deliberations on affective polarization is the ambitious field experiment by Fishkin et al. (2021). They brought together a representative sample of American voters to discuss five major political issues and demonstrate that cross-partisan deliberation reduces both policy and affective polarization. However, since all participants were in mixed-partisan groups, we cannot compare the effect of the echo chamber with cross-partisan dialogue. Yet observational data suggest that there is a strong association between the political homogeneity of friendship networks and affective polarization. For example, Figure 1 uses data from an original representative survey of British adults in July 2021² to show the relationship

² The survey was designed by the authors and conducted by well-established polling company YouGov. The sample is representative

FIGURE 1. How Affective Polarization Is Correlated with Perceptions of Friends' Voting Behavior

Note: 83% CI (thick), 95% CI (thin). Affective polarization is measured as the thermometer difference between parties.

between the standard measure of affective polarization (the thermometer score for someone's in-group party minus the thermometer score for their out-group party) and people's perceptions of how their friends voted in the 2019 U.K. general election.³

As Figure 1 shows, perceptions of the politics of one's friends are correlated with affective polarization. Conservative partisans who think "almost all" their friends voted Conservative score over 20 points higher on the affective polarization scale than Conservative partisans who think their friends are an equal mixture of Conservative and Labour voters. The same applies to Labour partisans. These are correlations, however. While it is plausible that

associating with fellow partisans increases affective polarization, it is also plausible that people who are affectively polarized choose friends who are similar to themselves. It is also possible that people from the other side do not want to be friends with people who are strongly affectively polarized. Given this, it is not surprising that Iyengar et al. (2019, 136) conclude in their review of the literature on affective polarization: "All told, therefore, it is premature to reach any firm conclusions about the role of 'echo chambers' either in person or online, as causes of affective polarization."

Since that conclusion, however, two important large-scale experimental studies have examined the causal effect of cross-party political discussion on affective polarization and found that it can reduce partisan animosity. Levendusky and Stecula (2021) ask participants to discuss different news articles about partisan polarization (given to the homogeneous partisan group) or partisan consensus (given to the heterogeneous partisan group). They find that the compound treatment of being in a cross-party group and reading about partisan consensus reduces affective polarization. In another large-scale experimental study, Santoro and Broockman (2022) isolate the effect of conversations between out-partisans (compared to co-partisans) and similarly find that video conversations between two out-partisans about "their perfect day" decrease affective polarization;

of the British population on key socioeconomic demographics and vote choice.

³ Perceptions of friends were measured using the question: "Overall, how do you think most of your closest friends voted in the 2019 General Election?" Response options were: Almost all voted Conservative, Most vote Conservative, An equal mixture voted Conservative and Labour, Mostly voted Labour, Almost all voted Labour, Mostly voted for a different party, Mostly didn't vote, Mostly too young to vote, and Don't know. Very few Conservative partisans think "almost all" their friends vote Labour and very few Labour partisans think "almost all" their friends vote Conservative so they are not shown. People living in Scotland are excluded due to the different party system.

however, when out-partisans discuss a topic that could engender partisan disagreement (in-party strengths/out-party flaws), they find no effects on affective polarization. Both studies break new ground by explicitly studying how the partisan heterogeneity of discussions can reduce affective polarization, but both only find effects when the discussion topic is not related to policy and is consensus-enhancing. One important question therefore, is whether these findings replicate the real-world echo chamber in which people often discuss actual policy issues and in which their policy views are echoed back at them or not. In this article, we thus examine whether discussion of a divisive and salient political issue within partisan echo chambers causes affective polarization when compared to discussion of the same issue in heterogeneous (mixed partisan) groups.

Our expectation is that discussions in homogeneous partisan groups influence not only individuals' policy opinions but also reinforce affective polarization. Two of the same mechanisms which likely underpin the echo chamber effect for policy polarization—social conformity and intergroup contact—also provide a rationale for how echo chambers impact affective polarization. First, homogeneous partisan networks can lead to stronger partisan identities through a mechanism of social conformity. People surrounded by just co-partisans will want to express their allegiance to that group and avoid the interpersonal cost associated with being seen as an outsider to the group (Schachter 1951). Second, intergroup contact theory teaches us that the composition of a group can shape how people view the out-group. According to Allport (1954), interaction within more heterogeneous networks can reduce group separation and prejudice. This is a form of perspective taking, where the presence of another partisan group allows individuals to see the other side as fellow humans with similar hopes and fears. Interaction with members of an out-group thus enables people to experience how out-group members view the world which reduces out-group animosity (Allport 1954; Pettigrew et al. 2011; Pettigrew and Tropp 2006). In contrast, if partisans on different sides are more isolated from each other and interact less in their real, or virtual lives, they are more likely to view the out-group negatively.

The fact that affective polarization is rooted in strong in-group attachment and out-group animosity leads to our expectations about the impact of the echo chamber on affective polarization, positive partisanship, and negative partisanship. We thus formulate the following hypotheses:

H2: *Homogeneous discussion groups will increase affective polarization compared to heterogeneous discussion groups.*

H2a: *Homogeneous discussion groups will increase positive partisanship compared to heterogeneous discussion groups.*

H2b: *Homogeneous discussion groups will increase negative partisanship compared to heterogeneous discussion groups.*

DATA AND METHODS

Experimental Design

We conduct a large-scale, pre-registered⁴ online experiment to test our hypotheses.⁵ This experiment makes a unique empirical contribution to the literature on affective polarization and echo chambers by adopting several novel design features. First, most experimental studies on group composition and polarization focus on providing a moderated deliberative discussion setting (Baron et al. 2021; Fishkin et al. 2021; Groenlund, Herne, and Setälä 2015; Hartman et al. 2022; Levendusky and Stecula 2021). This work has made an important contribution to how we understand the effectiveness of deliberative interventions in reducing polarization. However, in contrast to studies that seek to promote balanced deliberation through highly moderated discussions and impartial information provision, our experiment is intentionally designed with minimal intervention from the facilitators during the discussion. While the setting is obviously more structured than casual everyday encounters, the aim is to get closer to the type of political discussion that people typically experience: a discussion without thick deliberative rules of engagement and about a highly divisive political issue. Second, while early studies on echo chambers and polarization were often conducted as small lab experiments with convenience samples of university students, our experiment uses a large, diverse sample of people who are broadly similar to our population of interest (partisans). This is comparable to the larger samples which Deliberative Polling typically employs (Fishkin et al. 2021; Luskin, Fishkin, and Jowell 2002). Finally, while most previous studies on affective polarization focus on the United States, we build on recent work that expands this geographical focus to other Western democracies (Gidron, Adams, and Horne 2020; Hartveld 2021; Hobolt, Leeper, and Tilley 2021; Reiljan 2020; Wagner 2021). Specifically, we conduct our experiment in a party system which tends to show neither unusually high nor low levels of affective polarization: Britain. This provides a broadly representative case for testing the effects of group composition on polarization.

Our experiment consisted of 86 small group discussions with 724 participants that took place from September 20 to December 10, 2021. All participants

⁴ Pre-registration plan is here: <https://osf.io/29zb7>. Replication data and files, codebooks, and questionnaires can be found here: <https://doi.org/10.7910/DVN/BM2A1Q> (Hobolt, Lawall, and Tilley 2023). In this article, we test H3, H4, and H5 of the pre-registration plan. H1 and H6 relate to the emotiveness of the actual discussions and form a separate paper. H2 refers to the video affect level, which we varied from low emotional affect to high emotional affect with two different videos (see Appendix 3 of the Supplementary Material for more details).

⁵ Ethics approval was granted by the Research Ethics Committee of the London School of Economics and Political Science (No. 25671) and the Ethics Committee of the Centre for Experimental Social Sciences at the University of Oxford (No. OE_0061). See Hobolt, Lawall, and Tilley (2023).

TABLE 1. Overview of Treatment Groups

| Condition | Conservative | Labour | Total |
|----------------------|--------------|--------|-------|
| Mixed group | 156 | 155 | 311 |
| Homogeneous group | 127 | 220 | 347 |
| Non-discussion group | 20 | 46 | 66 |
| Total | 303 | 421 | 724 |

were Labour and Conservative partisans who lived in England.⁶ Potential participants took a pre-treatment survey and were then invited to sign up for an online discussion session. During the online session, held on Zoom, participants watched a very short introduction video about a specific immigration policy and then either discussed that policy in small groups of 6–8 people or went straight to a post-treatment survey. The policy discussed was “The New Plan for Immigration”: a set of government policies to change the asylum and immigration system, which was in the process of going through the legislature.

We use a 2×2 factorial design: the first factor is discussion group partisan composition (mixed partisan group/homogeneous partisan group) and the second factor is elite affective polarization (high elite affective polarization/low elite affective polarization). We designed the experiment so that the factors would be fully crossed, meaning that half of the participants in our study would be part of a mixed discussion group and half of the participants would be part of a homogeneous discussion group (see Appendix 2 of the Supplementary Material for full details of the experimental design). We focus exclusively on discussion group composition in this article, but see Appendix 3 of the Supplementary Material for further discussion of the elite polarization treatment. In terms of group composition, participants were assigned to either a homogeneous group with only co-partisans (e.g., Labour partisans were only placed with other Labour partisans), or to a heterogeneous group with a balanced mix of co-partisans and out-partisans (e.g., four Labour and four Conservative partisans in the group). Finally, a small number of participants attended a session and watched the video, but were not assigned to participate in a group discussion. Table 1 shows the numbers in each group (see also balance tables in Appendix 4 of the Supplementary Material).

Recruitment and Participants

To ensure high external validity, we recruited a large sample, broadly representative of the majority of the

⁶ We limit our focus to England in order to avoid the more complex party systems of Wales and Scotland. As of July 2021, according to our nationally representative survey, 58% of people in England identified as Conservative (33%) or Labour (24%) partisans. A further 9% identified with the Liberal Democrats and another 8% with other smaller parties (Greens, Reform UK, etc.).

British population who feel aligned with one of the two major parties. We recruited partisans using advertisements on Facebook. We chose Facebook because of its wide reach in Britain. In 2021, 65% of adults in Britain used Facebook, thereby giving us access to a more representative group of participants than traditional research lab subject pools (Ofcom 2021).⁷ Appendix 1 of the Supplementary Material discusses the recruitment process via Facebook in more detail.

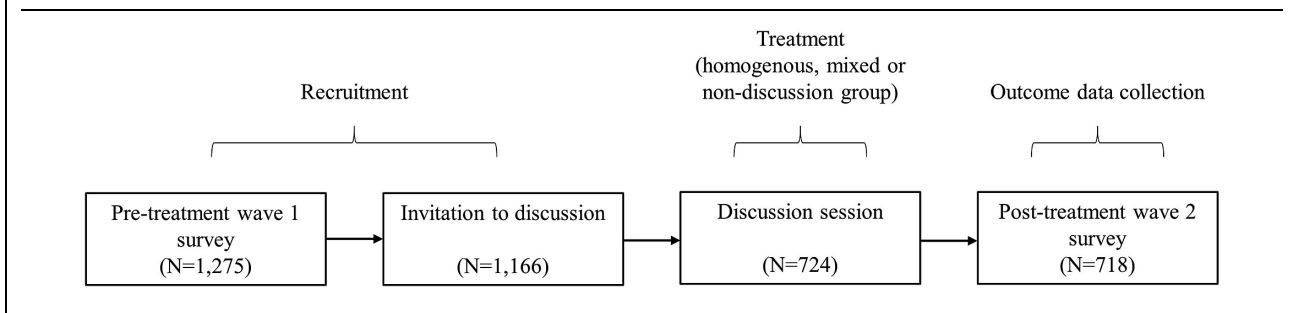
Figure 2 shows the flow of recruits into participants into completed Wave 2 surveys. Initial recruits were screened for partisanship of one of the two major parties (Conservative or Labour), age (over 18), nationality (British), and region (England), and then invited to fill in our Wave 1 survey. This asked questions on social demographics, political attitudes, attitudes toward the specific policy, measures of affective polarization, and specific measures of negative and positive partisanship. 1,275 people filled in the Wave 1 survey. The 1,166 respondents who identified as Labour or Conservative partisans, were over the age of 18, had British citizenship, lived in England, and agreed to the consent form were then invited via email to sign up to an online discussion event. A total of 724 participants attended an online discussion session, and, of those, almost all (718) filled in the Wave 2 survey. To incentivize participation, participants were paid £3 (3.50 USD) for completing the Wave 1 survey and £20 (24 USD) for taking part in the online session and filling in the Wave 2 survey.⁸

In Appendix 5 of the Supplementary Material, we compare the people who participated in the discussion sessions, and those who completed the pre-treatment survey, with a representative sample of the British population. To summarize, there are very few differences between the people who participated in discussions and those who took the survey but never participated in a discussion. Equally, our participants are similar demographically to partisans in the general population. For example, there are no large differences by age, race, or gender. Nonetheless, we should note that compared to the population, our sample is more educated, with almost twice as many holding a degree, and is two to three times as likely to discuss politics with friends regularly.⁹

⁷ It is worth noting that age differences between Facebook users and the general population are now also fairly small. According to our July 2021 representative survey, the 30 per cent of people who use Facebook less frequently than once a week in Britain have an average age of 51 (SD=17.6) whereas the 70 per cent of people who use Facebook weekly or more have an average age of 48 (SD=16.8).

⁸ We only allowed participants to take part in one discussion session, but we invited all participants who did not turn up to a discussion session to future sessions.

⁹ Fortunately, political engagement and education do not appear to interact with our treatment in predicting either policy or affective polarization. Appendix 6 of the Supplementary Material, which contains all the main models, also shows models that include interactions between our homogeneous/mixed group treatment and engagement/education. None of these interactions are statistically significant.

FIGURE 2. Experimental Design

Online Discussion Sessions and Group Treatment

Each of the online video discussion sessions lasted for around 45 minutes. After a 5-minute information video about a specific policy (the Conservative government's New Plan for Immigration), most participants had a 30-minute facilitated small group discussion in a breakout room and then filled in a post-treatment questionnaire. A smaller number of participants went straight to the post-treatment questionnaire after the video. The breakout rooms each had 6–8 participants plus a trained facilitator.¹⁰ Facilitators started the group discussion by calling on each participant to briefly give their opinion on the New Plan for Immigration and then opened up the discussion for follow-up comments that mostly focused on the policy or the wider issue of asylum seekers and immigration. Facilitators were told to let the discussion develop as naturally as possible and intervene only to prompt further discussion.

Crucially, we manipulated the partisan discussion environment. Participants were randomly assigned to either a homogeneous partisan group (with co-partisans) or a mixed partisan group (with the same number of Labour and Conservative partisans). Across the whole experiment, about half of the discussion groups were homogeneous and half were heterogeneous. The only determinant of assignment to a discussion group was the alphabetical ordering of a participant's first name in a given session's attendance list. This is independent of the participant's potential outcomes and therefore constitutes a valid random assignment mechanism. Participants did not know that the groups were arranged in this manner in advance, nor were they given explicit information about the partisan affiliations of their fellow group members at any point.

Outcome Measures

As outcome variables, we measure policy polarization, affective polarization, positive partisanship, and negative partisanship. All are measured identically in the Wave 1 and Wave 2 surveys.

Our measure of policy polarization focuses on the specific issue under discussion, the New Plan for Immigration, and measures how much people agree with three elements of the policy on 0–10 scales running from very strongly against to very strongly in favour. We combine these three specific questions into an additive scale, where higher values denote stronger agreement with the New Plan for Immigration. We also asked a general question about overall agreement with the policy. The wordings are below:

One part of the policy is that asylum seekers who arrive legally will be treated differently to those who arrive illegally. The government will try to remove anyone who arrives illegally and who could have claimed asylum in another safe country like France. Where do you stand on that policy?

Another part of the policy is that people who arrive illegally, and successfully claim asylum, will find it more difficult to claim benefits than asylum seekers who arrived legally and will have limited rights about reuniting with their family in Britain. Where do you stand on that policy?

Another part of the policy is a faster appeals process for asylum seekers. This aims to speed up the removal of people who are not granted asylum by limiting their rights to appeal. Where do you stand on that policy?

Thinking about the plan as whole, where do you stand on the "New Plan for Immigration?"

We use a standard measure of affective polarization: the difference between two questions that ask people how they feel about the two parties on a 0–100 thermometer scale (Gidron, Adams, and Horne 2020; Reiljan 2020; Wagner 2021). This measure thus focuses on the parties as objects of affect. Specifically, we ask respondents:

We'd like you to rate how you feel towards the Conservative party and the Labour party on a scale from 1–100, which we call a "feeling thermometer." Ratings between 0 and 49 mean that you feel unfavourable and cold. Ratings between 51 and 100 mean that you feel favourable and warm. A rating of 50 means that you have no feelings one way or the other. How would you rate your feelings towards the [political party]?

¹⁰ The majority of the facilitators (five out of six) were external to the research team. We did not share the objective of the experiment with these five facilitators. They were all native English speakers.

TABLE 2. Manipulation Check: Did Participants Learn Anything?

| | Conservative partisans | | Labour partisans | |
|--|------------------------|--------|------------------|--------|
| | Wave 1 | Wave 2 | Wave 1 | Wave 2 |
| Heard nothing about NPI (before today) | 29% | 4% | 26% | 5% |
| Mean Conservative Party position on NPI (1–10) | 8.6 | 8.5 | 9.2 | 8.7 |
| Mean Labour Party position on NPI (1–10) | 2.9 | 2.2 | 3.3 | 2.1 |

TABLE 3. Manipulation Check: Did the Treatment Affect Perceptions of Other Participants?

| | Conservative partisans | | Labour partisans | |
|---|------------------------|-------------|-------------------|-------------|
| | Homogeneous group | Mixed group | Homogeneous group | Mixed group |
| Most people in my group vote the same way | 77% | 19% | 94% | 21% |
| I agreed with most people in my group | 84% | 37% | 95% | 39% |

In principle, the difference between the two thus runs from -100 to $+100$, although in practice since extremely few people rate the other party as better than their own party, it runs from 0 to 100 with 100 as the maximum level of affective polarization.

To separately measure people's positive partisanship, we use a battery of five questions which focus on both parties and partisans as objects of affect. Variations of these questions have been used in the United States (Green, Palmquist, and Schickler 2004; Greene 2000; Huddy, Mason, and Aaroe 2015; Luttig 2018) and we replicate the battery used by Hobolt, Leeper, and Tilley (2021) to measure party identity in Britain. We ask people whether they agree or disagree with the following statements with regard to their own party and fellow partisans. The response options were strongly disagree to strongly agree which are scored 1–5 and then averaged.

When I speak about the [in-group party], I usually say “we” instead of “they.”

When people criticize the [in-group party], it feels like a personal insult.

I have a lot in common with other supporters of the [in-group party].

When I meet someone who supports the [in-group party], I feel connected with this person.

When people praise the [in-group party], it makes me feel good.

We also follow Iyengar, Sood, and Lelkes (2012) and measure positive trait stereotyping. Specifically, we asked people how well they thought two positive characteristics (honesty and intelligence) and two negative characteristics (selfishness and hypocrisy) described people on their own side. Agreement is on a 1–5 scale from 1 (not at all well) to 5 (very well). We combine the four items, reversing the second two responses, to make an additive scale, running from 1 to 5, measuring positive perceptions of the respondent's in-group partisans.

To measure negative partisanship, we use a five-item battery based on Bankert (2021). These questions are similar to the measure of positive partisanship but refer to out-group animosity toward the rival party and rival partisans. Again, the response options were strongly disagree to strongly agree, which are scored 1–5 and then averaged. The list of items is as below.

When [out-group party] does well in opinion polls, my day is ruined.

When people criticize the [out-group party], it makes me feel good.

I do not have much in common with [out-group party] supporters.

When I meet someone who supports [out-group party], I feel disconnected.

I get angry when people praise [out-group party].

We also measure negative trait stereotyping of the out-group. Here, we combine the four items measuring perceived characteristics of the other side, reversing the first two responses about honesty and intelligence, to make an additive scale, running from 1 to 5, measuring negative perceptions of out-group partisans. A distinct advantage of using these in-group and out-group partisan batteries of questions—in addition to the thermometer ratings—is that they allow us to capture affect toward both parties and partisans (Druckman et al. 2022; Druckman and Levendusky 2019).

Manipulation Checks

We include two types of manipulation checks in our experiment. First, we test whether people learned anything about the policy issue (Table 2) and second, we test whether people realized that their discussion group was homogeneous or not with regard to partisanship (Table 3). As Table 2 shows, most people had heard of the New Plan for Immigration when asked in the Wave

TABLE 4. OLS Regression Models Predicting Changes in Policy Opinions by Treatment among Conservative Partisans

| | NPI illegal entry | NPI benefit cut | NPI appeals reduced | NPI scale | NPI self- placement |
|----------------------|-------------------|-----------------|---------------------|-----------|---------------------|
| Homogeneous group | 0.42 | 0.62* | 0.62* | 0.54* | 0.52* |
| Non-discussion group | 0.29 | -0.28 | -0.43 | -0.15 | 0.15 |
| Mixed group | - | - | - | - | - |
| Low affect | -0.06 | 0.10 | -0.05 | 0.00 | 0.01 |
| High affect | - | - | - | - | - |
| Intercept | -0.16 | -0.16 | 0.02 | -0.15 | 0.01 |
| <i>N</i> | 295 | 290 | 286 | 279 | 288 |

Full model results in Appendix 6 of the Supplementary Material. * $p < 0.050$.

TABLE 5. OLS Regression Models Predicting Changes in Policy Opinions by Treatment among Labour Partisans

| | NPI illegal entry | NPI benefit cut | NPI appeals reduced | NPI scale | NPI self- placement |
|----------------------|-------------------|-----------------|---------------------|-----------|---------------------|
| Homogeneous group | -0.42* | -0.77* | -0.80* | -0.65* | -0.53* |
| Non-discussion group | 0.75* | 0.26 | -0.17 | 0.24 | 0.24 |
| Mixed group | - | - | - | - | - |
| Low affect | -0.24 | -0.01 | 0.20 | 0.03 | 0.08 |
| High affect | - | - | - | - | - |
| Intercept | -0.76* | -0.32 | -0.00 | -0.41* | -0.54* |
| <i>N</i> | 403 | 409 | 394 | 380 | 411 |

Full model results in Appendix 6 of the Supplementary Material. * $p < 0.050$.

1 survey: only about a quarter of both Labour and Conservative partisans had heard nothing about it. The questions on the policy in the Wave 1 survey were remembered, however, as virtually nobody said that they had not heard about the plan in the Wave 2 survey. The video and discussion seem to have led to people updating their information on the parties to some extent as well. The Labour Party, in particular, is seen as more opposed to the plan in Wave 2 than in Wave 1. Overall, participants appear to view the two parties as more polarized on the issue after the video and discussion. Nonetheless, these changes are fairly slight and it does not appear that any aspect of the treatment radically changed people's perceptions of party policy positions.

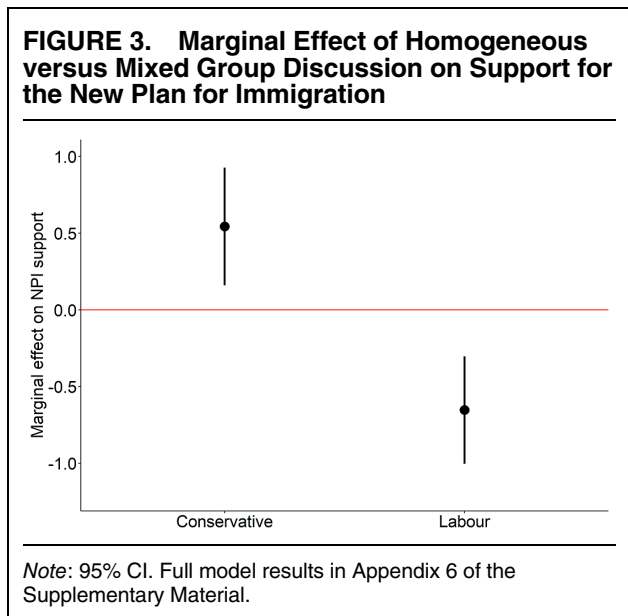
In terms of how people saw the discussion itself, Table 3 shows that people in homogeneous groups generally realized that they were in a politically homogeneous group. The percentages in the table refer to those who agreed or strongly agreed with the statements. While large majorities agreed that people in their group voted the same way in the homogeneous groups, few people thought that in the mixed groups. Although the differences are slightly less marked, people in homogeneous partisan groups also tended to find themselves in greater agreement with most people in their group than those in mixed groups.

RESULTS

Policy Polarization

To test our first hypothesis concerning the effect of group composition on policy polarization, Tables 4 and 5 show OLS regression models predicting the change from Wave 1 to Wave 2 on three different aspects of the New Plan for Immigration, an additive scale of those three questions and a self-placement scale of someone's overall attitude toward the plan. All measures run from 0 to 10 with 10 as the most supportive and 0 as the least supportive. Included as predictors of change are discussion type (homogeneous group, non-discussion group, and mixed group as the reference) and video affect (low and high as the reference).¹¹ We run these models separately for Labour and Conservative partisans since we expect the effects to run in opposite directions with Conservatives becoming more

¹¹ We also include, not shown, a measure of how late the discussion group occurred relative to other discussion groups. It was possible that attitudes toward the policy changed between early September and late December as it proceeded through the legislative process and inevitably people who participated in one of the later discussion groups were more likely to have completed the Wave 1 survey longer ago (the correlation between time elapsed between survey waves and date of discussion session is 0.35).



favorable toward the policy and Labour partisans becoming less favorable. By looking at within-person change, we hold constant all factors about individuals and can directly test whether our treatment changed people's minds.¹²

There is a straightforward pattern. Conservative partisans in the homogeneous groups became more supportive of the policy, however measured, when compared to Conservatives in mixed groups. Equally, Labour partisans in the homogeneous discussion groups became less supportive of the policy, however measured, when compared to those in mixed groups. This is demonstrated most clearly in Figure 3, which shows the marginal effect on policy support (the three-item scale) of discussion in a group of like-minded partisans compared to discussion in a mixed group by Labour and Conservative partisans.

In other words, homogeneous discussion groups produced policy polarization compared to mixed discussion groups: the echo chamber caused policy polarization compared to a mixed social environment. This is despite the fact that Conservative and Labour partisans are already highly polarized on the issue of immigration. Whereas Conservative partisans who participated in the discussion place themselves at 7.8 on average in the Wave 1 survey, Labour partisans place themselves at 2.2. Even though there might appear little room to do so, we increase the partisan gap by well over one point on the 0–10 scale for those in homogeneous groups compared to those in mixed groups. This provides strong support for Hypothesis 1 that echo chambers polarize people's policy attitudes (compared to mixed discussion settings). The fact that this happens in the course of a 30-minute discussion with strangers over Zoom on an issue that is already

¹² In line with our pre-registered analysis strategy, we also show additional models in Appendix 6 of the Supplementary Material with raw post-treatment means as the outcome.

highly divisive by partisanship, shows the potential importance of homogeneous partisan networks in creating politically polarized societies.

Affective Polarization

Next, we test Hypothesis 2: whether homogeneous groups increase affective polarization compared to heterogeneous groups. As discussed, our main measure of affective polarization is the commonly used scale which runs from –100 to +100. The value +100 in this case indicates that someone gives a score of 100 on a 0–100 liking thermometer to their own party and 0 to the other party; conversely –100 means that someone gave 100 to the other party and 0 to their own party. The left-hand panel in Figure 4 shows the pre-treatment and post-treatment scores on this scale for the homogeneous and mixed groups.

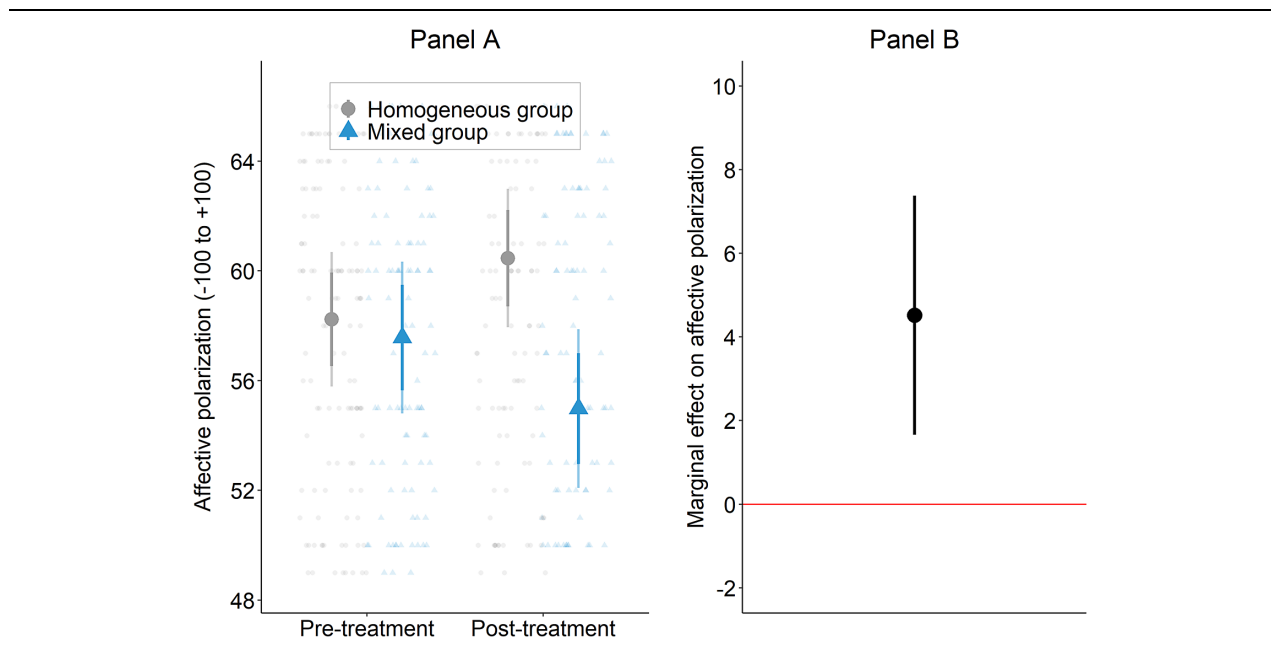
The right-hand panel of Figure 4 shows the marginal effect, estimated from an OLS model (full results in Appendix 6 of the Supplementary Material), of being assigned to a homogeneous group rather than a mixed group on the difference in affective polarization scores between Wave 2 and Wave 1. This marginal effect is about 5 points, or a little less than a quarter of a standard deviation, in affective polarization. Someone who talked about immigration policy with a group of like-minded partisans, rather than in a mixed group of both in-partisans and out-partisans, scores 5 points higher on the thermometer scale of affective polarization.¹³ In Appendix 6 of the Supplementary Material, we also show that these results hold when we estimate treatment effects on raw post-treatment outcomes rather than the change from Wave 1 to Wave 2.

Additionally, we look at measures of positive and negative partisanship (H2a and H2b). Figure 5 shows the marginal effects of being in a homogeneous versus mixed group from OLS models predicting the difference in scores between Wave 2 and Wave 1 on our four 1–5 scales of positive and negative partisan affect. Positive coefficients indicate greater positive partisanship and greater negative partisanship—that is, greater affective polarization. All the effects are positive, so homogeneous groups become more affectively polarized than mixed groups on these four measures, as expected. However, only the effects for positive partisanship and negative out-group trait perceptions are statistically significant at the 5% level.

We can also break down our results for affective polarization by the party identity of our respondents. Figure 6 shows the marginal treatment effects for the thermometer measure of affective polarization separately for Labour and Conservative partisans. As

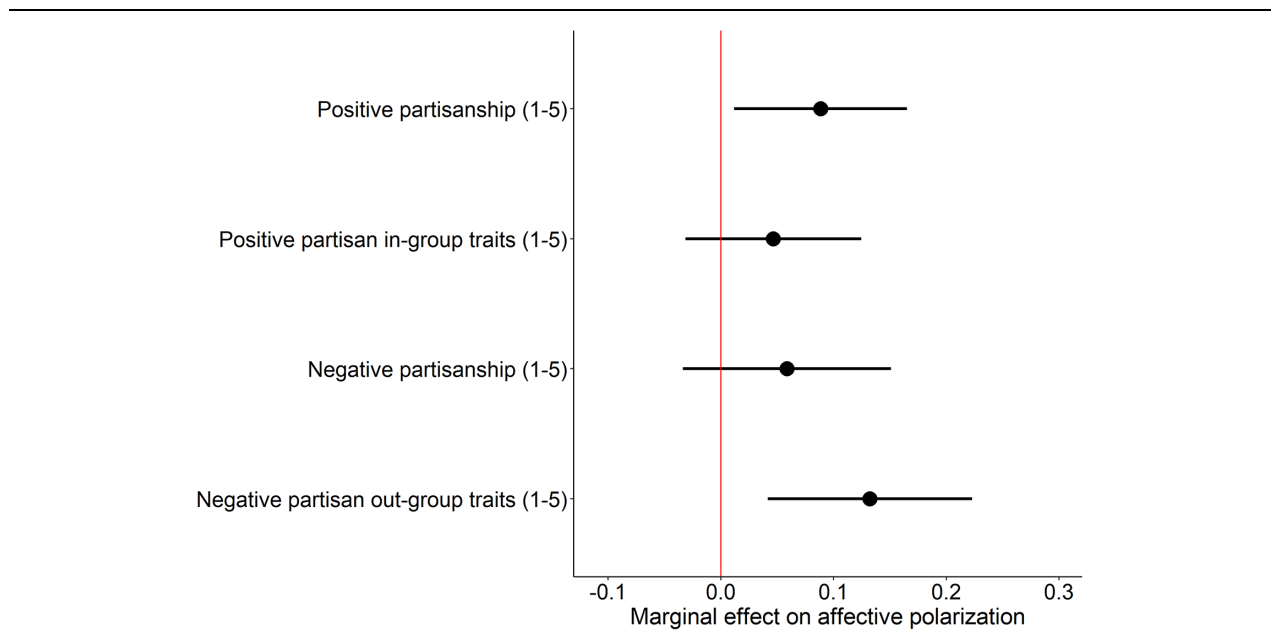
¹³ Participants in homogeneous partisan groups thus became significantly more affectively polarized after discussion (paired *t*-test mean change = 1.97, $p < 0.05$; see Appendix 6 of the Supplementary Material), while mixed groups become significantly less affectively polarized after discussion (paired *t*-test mean change = –3.2, $p < 0.01$; see Appendix 6 of the Supplementary Material).

FIGURE 4. The Effect of Group Composition on Affective Polarization



Note: Panel A: Mean affective polarization levels by experimental condition and time point. 83% CI (thick), 95% CI (thin). Panel B: Marginal effect of homogeneous versus mixed group on change in affective polarization. 95% CI. Excludes respondents who changed partisanship from W1 to W2 ($N = 21$). Full model results in Appendix 6 of the Supplementary Material.

FIGURE 5. Marginal Effect of Homogeneous versus Mixed Group on Affective Polarization

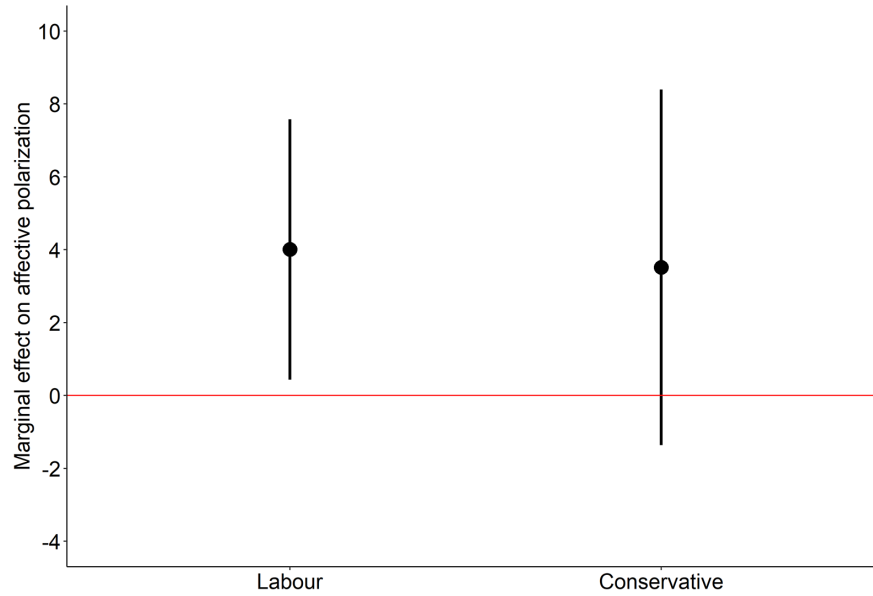


Note: 95% CI. Excludes respondents who changed partisanship from W1 to W2 ($N = 21$). Full model results in Appendix 6 of the Supplementary Material.

Figure 6 shows, the size of the effects on the thermometer scale are very slightly larger for Labour supporters, and the effect for Conservatives is not statistically significant at conventional significance levels. Figure 7 shows

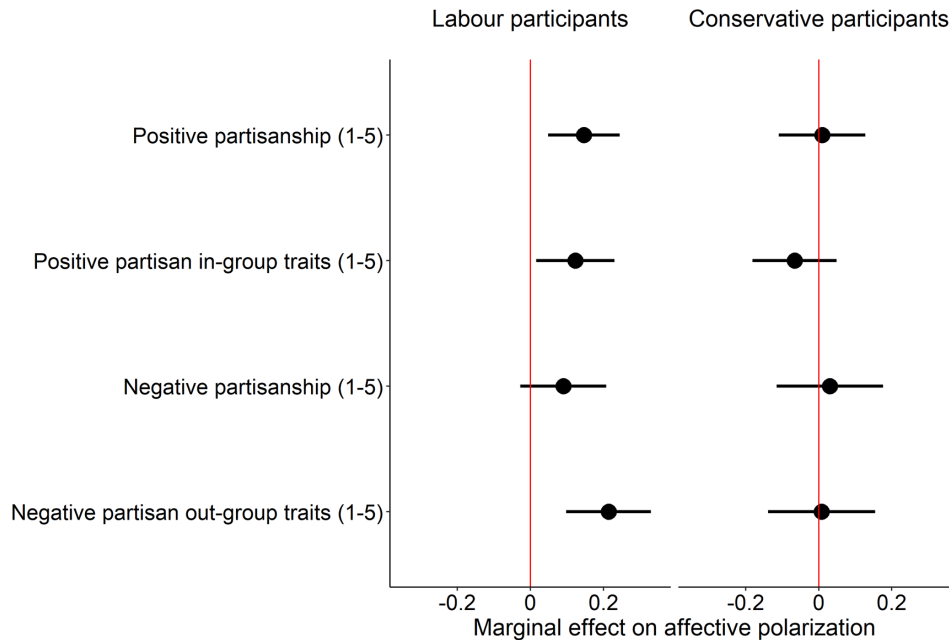
the treatment effects for Labour and Conservative partisans for the positive and negative partisanship outcomes. The treatment effects we see in Figure 5 are essentially driven by Labour supporters: the difference

FIGURE 6. Marginal Effect of Homogeneous versus Mixed Group on Affective Polarization by Partisanship



Note: 95% CI. Excludes respondents who changed partisanship from W1 to W2 ($N = 21$). Full model results in Appendix 6 of the Supplementary Material.

FIGURE 7. Marginal Effect of Homogeneous versus Mixed Group on Affective Polarization by Partisanship



Note: 95% CI. Excludes respondents who changed partisanship from W1 to W2 ($N = 21$). Full model results in Appendix 6 of the Supplementary Material.

between Labour partisans in homogeneous groups and mixed groups grows substantially after treatment. The picture for Conservative partisans is different. Here,

there is no evidence that our treatments affected the scale scores of positive and negative partisanship. Overall, while we thus find support for our pre-registered

hypotheses—homogeneous groups are more positive toward people on their own side and more negative toward people on the other side than heterogeneous groups—this effect is only consistently, and statistically significantly, apparent for Labour partisans.

This difference between Labour and Conservatives, which we do not find for policy polarization, is not straightforward to explain. One possibility is that there is something about Conservative and Labour partisans, whether in reality or in our sample, that differs and then alters how they respond to the treatment with regard to affective polarization. There are two obvious factors: starting levels of partisan strength and deep-seated personality traits. Both may affect reactions to social situations such as a political discussion. We can test these directly as we have measures of partisan strength at Wave 1 and we also measured the Big Five personality traits, as well as narcissism, at Wave 1. It is certainly the case that Labour and Conservative partisans differ on these dimensions. Conservatives are less strongly identified with their party than Labour supporters at Wave 1; there are also substantial differences between the two sides in several personality traits. In neither case, however, do we find any moderating effects. Weak and strong partisans react similarly to the treatment and people who differ on the Big Five traits and narcissism also react in a similar manner. We show the (null) results of these models in Appendix 7 of the Supplementary Material.

The other main possibility with which we are left is that incumbency and opposition dynamics are important. After all, Labour partisans were attached to a party which was in opposition and Conservative partisans were attached to the ruling party.¹⁴ Incumbency matters because it may influence how different groups of partisans perceive their relative group status. A central tenet of social identity theory is that individuals adopt identities to develop a positive self-image (Tajfel and Turner 1979). If a party is in opposition, it is a “loser” which challenges how much self-worth opposition partisans can derive from their partisan identity and encourages out-group derogation. Partisans of the governing party, on the other hand, may feel less pressure to reaffirm their in-group identity because their party is a winner. Incumbency status may thus shield government partisans from the more emotive and affective reactions to network composition. Clearly, this is speculation, and we can only offer it as a tentative explanation for the differences we find between Conservative and Labour partisans. Nonetheless, a potentially fruitful avenue for future research would be to test explicitly whether incumbency and opposition status shape partisans’ affective responses to echo chambers.

¹⁴ Opposition-government dynamics may also have been particularly salient during our experimental treatment since the discussion was about a government immigration bill and the video showed clips from the House of Commons in which government ministers introduced the policy.

CONCLUSION

Given the heightened partisan divisions that we see in many democracies around the world, it is crucial to understand what drives political polarization. One concern is that partisans on different sides are increasingly isolated from each other and interact less in their real, or virtual lives, and this exacerbates partisan polarization. Our lab-in-the-field experiment shows that echo chambers do indeed increase both policy and affective polarization. This experiment is innovative in several ways. The key treatment—the partisan composition of the groups as either partisan echo chambers or mixed groups—enables us to establish the causal effect of the homogeneity of a discussion group on affective polarization. Our sample is a large representative group of British Labour and Conservative partisans who discuss a salient and divisive political issue, immigration policy, in a setting without strict deliberative norms or guidance. This combination allows us to test causally the effect of the inter-group context on both policy and affective polarization.

Consistent with earlier findings (most notably Klar 2014), we show that discussion with people from the same side, compared to discussion with a mixed group of people, increased policy polarization. This is even though Labour and Conservative partisans were already highly polarized on the issue of immigration before entering the discussion. More importantly, we also show that group discussion in the echo chamber raised affective polarization. People in homogeneous, compared to heterogeneous, discussion groups became more positive toward their own party and fellow partisans, but also more negatively disposed toward the other party and rival partisans. These findings thus provide important evidence that interpersonal networks—whether in real life or on social media—can play a role in shaping affective polarization. Our results are consistent with recent experiments conducted in the United States that highlight the depolarizing potential of cross-partisan discussions (Fishkin et al. 2021; Levendusky and Stecula 2021; Santoro and Broockman 2022), but they demonstrate that such effects can be found even in settings where partisans are encouraged to discuss divisive policy issues and with minimal moderation of the discussion.

There are, of course, limitations to our study. Some relate to the nature of studying group discussions experimentally. While this brings important methodological advantages, it also raises questions about whether the findings can be generalized to the type of real-life political discussions that may happen in the pub after work. We designed our experiment to make the discussions resemble, as far as possible, real face-to-face discussions in an online setting that had become common during the pandemic. While our participants are broadly representative of partisans in England, they do have higher levels of political interest than even the average partisan. Arguably, however, it is such politically interested partisans who are key opinion-formers and thus shape polarization in the electorate. Another possible limitation is that our experiment focused on a

single policy issue, and this raises questions about whether we would find similar effects on other issues. We cannot answer this, but since Santoro and Brookman (2022) find stronger effects of cross-partisan discussion on affective polarization when opposing partisans discuss issues that do not give cause for disagreement, it seems likely that our findings may generalize to issues that are less divisive.¹⁵ Finally, our experiment was conducted in the United Kingdom, which, unlike most established democracies, has a dominant two-party system. This feature makes it more straightforward to allocate people to homogeneous and mixed groups, but the dynamics of how in-group and out-group partisans relate to each other may be more complicated in other countries. Hopefully, future studies can explore these group dynamics in other settings.

What are the wider implications of our findings? On the one hand, we might think that the glass is half empty. Increased policy and affective polarization among those who discuss politics with like-minded people is concerning since there is evidence that people increasingly find themselves in partisan echo chambers, both in real life and on social media (Mason 2018; Sunstein 2018). Given the potential consequences of affective polarization on the quality of democratic deliberation and support for democratic norms, it is worrying that out-group animosity, in particular, is enhanced by homogeneous partisan networks, and that this appears to be driven in large part by opposition partisans.

On the other hand, we might think the glass is half full. One implication of our results is that reducing the degree to which people find themselves in homogeneous echo chambers will also reduce polarization. Indeed, people appear to become less affectively polarized after discussing politics with a mixed group of partisans. This is a remarkable finding, given that our experiment is not set up in a way to encourage strong deliberative norms and cooperation. Instead, our experiment tests whether inter-group contact works even under realistic, and therefore suboptimal, conditions: a highly polarizing discussion task that does not encourage cooperation and a relatively short non-repeated interaction that does not allow people to establish lasting non-hierarchical social relationships (Allport 1954; Pettigrew et al. 2011). All of this suggests that institutions and interventions that encourage more political discussion, however limited, across the partisan divide may be successful in limiting out-group animosity.

SUPPLEMENTARY MATERIAL

The supplementary material for this article can be found at <https://doi.org/10.1017/S0003055423001211>.

¹⁵ See also Fishkin et al. (2021), who find a depolarizing effect of mixed group discussion across five different issues (including immigration).

DATA AVAILABILITY STATEMENT

Research documentation and data that support the findings of this study are openly available at the American Political Science Review Dataverse: <https://doi.org/10.7910/DVN/BM2A1Q>.

ACKNOWLEDGMENTS

We are grateful to the Nuffield College Centre for Experimental Social Science (CESS), especially Noah Bacine, Ray Duch, and Alessandro de Oliveira, and we thank Max Afnan, Chris Bick, Ceren Cinar, Andrew Lewis, Nick Lewis, and Ariel Perkins for their excellent research assistance. Previous versions of this article have been presented at MPSA 2022, EPSA 2022, Bocconi University, Durham University, Essex University, Konstanz University, LSE, Royal Holloway, University of Southern Denmark, and University of Zurich, and we are grateful to participants for their helpful comments. Finally, we would also like to thank the Editors and reviewers of the *APSR* for their insightful suggestions.

FUNDING STATEMENT

This research was generously supported by the Economic and Social Research Council (Grant No. ES/V004360/1; PIs: Sara B. Hobolt and James Tilley) and the European Research Council (Grant No. ERC-CoG-2014 647835-EUDEMOS; PI: Sara B. Hobolt).

CONFLICT OF INTEREST

The authors declare no ethical issues or conflicts of interest in this research.

ETHICAL STANDARDS

The authors declare the human subjects research in this article was reviewed and approved by the Research Ethics Committee of the London School of Economics and Political Science (no. 25671) and the Ethics Committee of the Centre for Experimental Social Sciences at the University of Oxford (no. OE 0061). The authors affirm that this article adheres to the APSA's Principles and Guidance on Human Subject Research.

REFERENCES

- Allport, Gordon W. 1954. *The Nature of Prejudice*. Cambridge: Addison-Wesley.
- Banda, Kevin K., and John Cluverius. 2018. "Elite Polarization, Party Extremity, and Affective Polarization." *Electoral Studies* 56: 90–101.
- Bankert, Alexa. 2021. "Negative and Positive Partisanship in the 2016 US Presidential Elections." *Political Behavior* 43 (4): 1467–85.

- Barbera, Pablo. 2020. "Social Media, Echo Chambers, and Political Polarization." In *Social Media and Democracy*, eds. Nathaniel Persily and Joshua A. Tucker, 34–55. Cambridge: Cambridge University Press.
- Barbera, Pablo, John T. Jost, Jonathan Nagler, Joshua A. Tucker, and Richard Bonneau. 2015. "Tweeting from Left to Right: Is Online Political Communication More Than an Echo Chamber?" *Psychological Science* 26 (10): 1531–42.
- Baron, Hannah, Robert Blair, Donghyun D. Choi, Laura Gamboa, Jessica Gottlieb, Amanda Lea Robinson, Steven Rosenzweig, et al. 2021. "Can Americans Depolarize? Assessing the Effects of Reciprocal Group Reflection on Partisan Polarization." Unpublished Manuscript. <https://doi.org/10.31219/osf.io/3x7z8>
- Broockman, David, Joshua Kalla, and Sean Westwood. 2023. "Does Affective Polarization Undermine Democratic Norms or Accountability? Maybe Not." *American Journal of Political Science* 67 (3): 808–28.
- Druckman, James N., Samara Klar, Yanna Krupnikov, Matthew Levendusky, and John Barry Ryan. 2022. "(Mis) Estimating Affective Polarization." *Journal of Politics* 84 (2): 1106–17.
- Druckman, James N., and Matthew S. Levendusky. 2019. "What Do We Measure When We Measure Affective Polarization?" *Public Opinion Quarterly* 83 (1): 114–22.
- Esterling, Kevin M., Archon Fung, and Taeku Lee. 2021. "When Deliberation Produces Persuasion Rather Than Polarization: Measuring and Modelling Small Group Dynamics in a Field Experiment." *British Journal of Political Science* 51 (2): 666–84.
- Festinger, Leon. 1950. "Informal Social Communication." *Psychological Review* 57 (5): 271–82.
- Fishkin, James, Alice Siu, Larry Diamond, and Norman Bradburn. 2021. "Is Deliberation an Antidote to Extreme Partisan Polarization? Reflections on 'America in One Room.'" *American Political Science Review* 115 (4): 1464–81.
- Gentzkow, Matthew, and Jesse M. Shapiro. 2011. "Ideological Segregation Online and Offline." *Quarterly Journal of Economics* 126 (4): 1799–839.
- Gidron, Noam, James Adams, and Will Horne. 2020. *American Affective Polarization in Comparative Perspective*. Cambridge: Cambridge University Press.
- Green, Donald P., Bradley Palmquist, and Eric Schickler. 2004. *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. New Haven, CT: Yale University Press.
- Greene, Steven. 2000. "The Psychological Sources of Partisan-Leaning Independence." *American Politics Quarterly* 28 (4): 511–37.
- Groenlund, Kimmo, Kaisa Herne, and Maija Setälä. 2015. "Does Enclave Deliberation Polarize Opinions?" *Political Behavior* 37: 995–1020.
- Harteveld, Eelco. 2021. "Fragmented Foes: Affective Polarization in the Multiparty Context of the Netherlands." *Electoral Studies* 71: 102332. <https://doi.org/10.1016/j.electstud.2021.102332>
- Hartman, Rachel, Will Blakey, Jack Womick, Chris Bail, Eli J. Finkel, Hahrie Han, John Sarrouf, et al. 2022. "Interventions to Reduce Partisan Animosity." *Nature Human Behaviour* 6 (9): 1194–205.
- Hetherington, Marc J. 2009. "Putting Polarization in Perspective." *British Journal of Political Science* 39 (2): 413–48.
- Hetherington, Marc J., and Thomas J. Rudolph. 2015. *Why Washington Won't Work: Polarization, Political Trust, and the Governing Crisis*. Chicago, IL: University of Chicago Press.
- Hobolt, Sara B., Katharina Lawall, and James Tilley. 2023. "Replication Data for: The Polarizing Effect of Partisan Echo Chambers." Harvard Dataverse. Dataset. <https://doi.org/10.7910/DVN/BM2A1Q>
- Hobolt, Sara B., Thomas Leeper, and James Tilley. 2021. "Divided by the Vote: Affective Polarization in the Wake of the Brexit Referendum." *British Journal of Political Science* 51 (4): 1476–93.
- Huber, Gregory A., and Neil Malhotra. 2017. "Political Homophily in Social Relationships: Evidence from Online Dating Behavior." *Journal of Politics* 79 (1): 269–83.
- Huddy, Leonie, Lilliana Mason, and Lene Aaroe. 2015. "Expressive Partisanship: Campaign Involvement, Political Emotion, and Partisan Identity." *American Political Science Review* 109 (1): 1–17.
- Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22: 129–46.
- Iyengar, Shanto, Guarav Sood, and Yphtach Lelkes. 2012. "Affect, Not Ideology: A Social Identity Perspective on Polarization." *Public Opinion Quarterly* 76 (3): 405–31.
- Iyengar, Shanto, and Sean J. Westwood. 2015. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." *American Journal of Political Science* 59 (3): 690–707.
- Kekkonen, Arto, and Tuomas Yl-Anttila. 2021. "Affective Blocs: Understanding Affective Polarization in Multiparty Systems." *Electoral Studies* 72: 102367. <https://doi.org/10.1016/j.electstud.2021.102367>
- Kingzette, Jon, James N. Druckman, Samara Klar, Yanna Krupnikov, Matthew Levendusky, and John Barry Ryan. 2021. "How Affective Polarization Undermines Support for Democratic Norms." *Public Opinion Quarterly* 85 (2): 663–77.
- Klar, Samara. 2014. "Partisanship in a Social Setting." *American Journal of Political Science* 58 (3): 687–704.
- Layman, Geoff C., Thomas M. Carsey, and Juliana Menasce Horowitz. 2006. "Party Polarization in American Politics: Characteristics, Causes, and Consequences." *Annual Review of Political Science* 9: 83–110.
- Lelkes, Yphtach, Gaurav Sood, and Shanto Iyengar. 2017. "The Hostile Audience: The Effect of Access to Broadband Internet on Partisan Affect." *American Journal of Political Science* 61 (1): 5–20.
- Levendusky, Matthew. 2013. "Partisan Media Exposure and Attitudes Toward the Opposition." *Political Communication* 30 (4): 565–81.
- Levendusky, Matthew S., James N. Druckman, and Audrey McLain. 2016. "How Group Discussions Create Strong Attitudes and Strong Partisans." *Research and Politics* 3 (2): 1–6.
- Levendusky, Matthew S., and Dominik A. Stecula. 2021. *We Need to Talk: How Cross-Party Dialogue Reduces Affective Polarization*. Cambridge: Cambridge University Press.
- Lipset, Seymour Martin. 1959. *Political Man: The Social Bases of Politics*. Baltimore, MD: Johns Hopkins University Press.
- Luskin, Robert C., James S. Fishkin, and Roger Jowell. 2002. "Considered Opinions: Deliberative Polling in Britain." *British Journal of Political Science* 32 (3): 455–87.
- Luttig, Matthew D. 2018. "The Prejudiced Personality and the Origins of Partisan Strength, Affective Polarization, and Partisan Sorting." *Political Psychology* 39 (1): 239–56.
- MacKuen, Michael, Jennifer Wolak, Luke Keele, and George E. Marcus. 2010. "Civic Engagements: Resolute Partisanship or Reflective Deliberation." *American Journal of Political Science* 54 (2): 440–58.
- Mason, Lilliana. 2015. "I Disrespectfully Agree: The Differential Effects of Partisan Sorting on Social and Issue Polarization." *American Journal of Political Science* 59 (1): 128–45.
- Mason, Lilliana. 2018. *Uncivil Agreement: How Politics Became Our Identity*. Chicago, IL: University of Chicago Press.
- McCoy, Jennifer, Tahmina Rahman, and Murat Somer. 2018. "Polarization and the Global Crisis of Democracy: Common Patterns, Dynamics, and Pernicious Consequences for Democratic Politics." *American Behavioral Scientist* 62 (1): 16–42.
- Muradova, L. 2021. "Seeing the Other Side? Perspective-Taking and Reflective Political Judgements in Interpersonal Deliberation." *Political Studies* 69 (3): 644–64.
- Mutz, Diana. 2006. *Hearing the Other Side: Deliberative versus Participatory Democracy*. New York: Cambridge University Press.
- Nyhan, Brendan, and Jason Reifler. 2010. "When Corrections Fail: The Persistence of Political Misperceptions." *Political Behavior* 32: 303–30.
- Nyhan, Brendan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y Chen, Hunt Allcott, et al. 2023. "Like-Minded Sources on Facebook Are Prevalent but not Polarizing." *Nature* 620 (7972): 137–44.
- Ofcom. 2021. "Adults' Online Behaviours and Attitudes 2021." Report.
- Pariser, Eli. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. London: Penguin.

- Pettigrew, Thomas F., and Linda R. Tropp. 2006. "A Meta-Analytic Test of Intergroup Contact Theory." *Journal of Personality and Social Psychology* 90 (5): 751–83.
- Pettigrew, Thomas F., Linda R. Tropp, Ulrich Wagner, and Oliver Christ. 2011. "Recent Advances in Intergroup Contact Theory." *International Journal of Intercultural Relations* 35 (3): 271–80.
- Reiljan, Andres. 2020. "Fear and Loathing across Party Lines (also) in Europe: Affective Polarization in European Party Systems." *European Journal of Political Research* 59 (2): 376–96.
- Rogowski, Jon C., and Joseph L. Sutherland. 2016. "How Ideology Fuels Affective Polarization." *Political Behavior* 38 (2): 485–508.
- Santoro, Erik, and David E. Broockman. 2022. "The Promise and Pitfalls of Cross-Partisan Conversations for Reducing Affective Polarization: Evidence from Randomized Experiments." *Science Advances* 8 (25): eabn5515.
- Schachter, Stanley. 1951. "Deviation, Rejection, and Communication." *Journal of Abnormal and Social Psychology* 46 (2): 190–207.
- Schkade, David, Cass R. Sunstein, and Reid Hastie. 2010. "When Deliberation Produces Extremism." *Critical Review* 22 (2–3): 227–52.
- Settle, Jaime E. 2018. *Frenemies: How Social Media Polarizes America*. Cambridge: Cambridge University Press.
- Strandberg, Kim, Staffan Himmelroos, and Kimmo Groenlund. 2019. "Do Discussions in Like-Minded Groups Necessarily Lead to More Extreme Opinions? Deliberative Democracy and Group Polarization." *International Political Science Review* 40 (1): 41–57.
- Suhay, Elizabeth, Emily Bello-Pardo, and Brianna Maurer. 2018. "The Polarizing Effects of Online Partisan Criticism: Evidence from Two Experiments." *International Journal of Press/Politics* 23 (1): 95–115.
- Sunstein, Cass R. 2009. *Going to Extremes: How Like Minds Unite and Divide*. Oxford: Oxford University Press.
- Sunstein, Cass R. 2018. *Republic: Divided Democracy in the Age of Social Media*. Princeton, NJ: Princeton University Press.
- Tajfel, Henri, and John C. Turner. 1979. "An Integrative Theory of Intergroup Conflict." In *Organizational Identity: A Reader*, eds. Mary Jo Hatch and Majken Schultz, 56–65. Oxford: Oxford University Press.
- Terren, Ludovic, and Rosa Borge-Bravo. 2021. "Echo Chambers on Social Media: A Systematic Review of the Literature." *Review of Communication Research* 9: 99–118.
- Toernberg, Petter, Claes Andersson, Kristian Lindgren, and Sven Banisch. 2021. "Modeling the Emergence of Affective Polarization in the Social Media Society." *PLoS One* 16 (10): e0258259. <https://doi.org/10.1371/journal.pone.0258259>
- Torcal, Mariano, and Josep M. Comellas. 2022. "Affective Polarisation in Times of Political Instability and Conflict. Spain from a Comparative Perspective." *South European Society and Politics* 27 (1): 1–26.
- Visser, Penny S., and Robert R. Mirabile. 2004. "Attitudes in the Social Context: The Impact of Social Network Composition on Individual-Level Attitude Strength." *Journal of Personality and Social Psychology* 87 (6): 779–95.
- Wagner, Markus. 2021. "Affective Polarization in Multiparty Systems." *Electoral Studies* 69: 102199. <https://doi.org/10.1016/j.electstud.2020.102199>