

Offline partisan segregation exceeds online echo chambers in the United States: Evidence from 1 million American voters

Megan A. Brown^{1,3}, Tiago Ventura^{2,3}, Joshua A. Tucker³, and Jonathan Nagler³

¹School of Information, University of Michigan

²McCourt School of Public Policy, Georgetown University

³Center for Social Media and Politics (CSMaP), New York University

Abstract

Social media is often blamed for the creation of echo chambers. However, these claims fail to consider the prevalence of *offline* echo chambers resulting from high levels of geographic partisan segregation in the United States. Our article empirically assesses these online versus offline dynamics by linking a novel dataset of voters' offline geographic partisan segregation extracted from publicly available voter files for 180 million US voters with their online network segregation on Twitter. We investigate offline and online partisan segregation using measures of geographical and online isolation of every matched voter-Twitter user to their co-partisans offline and online. Our results show that while social media users do tend to form politically homogeneous online networks, these levels of partisan sorting are significantly lower than those found in their geographical offline settings. Thus social media is providing users with networks that are more heterogeneous than their offline environments. We also find that Democrats are more ideologically isolated than Republicans in both settings. Our results contribute to the emerging literature on political communication and place the homophily of online networks in a broader context, providing novel evidence on partisan sorting both online and offline.

Introduction

The consolidation of social media and new digital technologies in today’s information environment has been widely cited in both popular (1; 2; 3) and academic debates (4; 5; 6; 7) as a critical driver of the recent rise in political polarization among the American electorate. The “echo chambers” narrative is often invoked as the primary mechanism by which social media causes polarization (4; 8; 9). In this narrative, social media has facilitated the formation of overly homogenous online communities—“echo chambers”—in which citizens more easily isolate themselves from people from opposing parties, and that hold distinct views about the world, consequently reinforcing their system of beliefs and reducing cross-cutting interactions with others who hold views different from the individual’s own views.

Yet, a growing body of research relying on a variety of data sources has provided a more nuanced view of this narrative, showing that the prevalence of social media echo chambers has been largely overstated (10; 11; 12; 13). Although these studies somewhat converge on the idea that online echo chambers may be an overstatement, the impact of social media on polarization remains a persistent puzzle; using digital field experiments in which authors manipulate components of users’ social media experience, or even incentivize users to deactivate their account for a certain period of time, some research shows that social media increases polarization (14; 15; 16), some shows that social media decreases polarization (17; 18), and other research shows no effect of social media on polarization (19; 20; 21; 22; 23; 24; 25).

To understand this puzzle, recent research has advanced an alternative causal mechanism in which social media effects must be considered conditional on users’ offline experiences and context (26; 18; 6). Relative to individuals’ offline environments, social media might actually increase exposure to outgroup users (18), encouraging them to engage

with content they would hardly encounter outside social media (6) and forcing them to interact beyond their local bubble (26). Because social media platforms often amplify extreme voices, and emotional and negative language (27; 28), this line of work shows that these interactions can often activate partisan identity, and lead to misperceptions about outgroup voters and uncivil online exchanges, which could increase polarization. A recent field experiment manipulating Twitter’s feed algorithm illustrates this dynamic (22). The authors find that assigning users to an algorithmic feed, compared to chronological ranking, increases users’ ties to more politically conservative accounts and shifts their views towards more conservative positions. Given the more liberal-leaning Twitter user base, as our data show, these algorithmic recommendations actually increase users’ exposure to outgroup content, rather than sorting users into echo chambers.

Our paper fills this gap by measuring how Americans’ online communities differ from, or resemble, their offline local communities. To answer this question, we develop novel individual-level measures of offline and online partisan sorting using a dataset of 950,000 Twitter¹ users matched to information on the users contained in augmented voter files. For offline partisan segregation, we use voter file information of our matched Twitter users to measure the physical proximity between every matched user and their co-partisan nearest neighbors. For online partisan segregation, we collect the Twitter networks of each matched user in our panel and estimate the ideology of each user that our matched voter follows using item-response estimates developed by (29). Across both measures, we consider voters to be segregated when the composition of their offline and online networks is highly homogenous, containing a high share of co-partisan voters/social media users. We refer to this measure of partisan segregation as online (offline) isolation, indicating that citizens in their online bubbles (local neighborhoods) are unlikely to interact with and be exposed to out-partisans

¹At the time of data collection, the platform was still called Twitter. We use Twitter throughout the paper.

Our study presents four core findings. First, we find that for both Republicans and Democrats, partisan segregation is *higher* in their offline environments than in their online environments. That is, even though their online environments contain more politically like-minded people than not, their offline environments are *even more* politically homogeneous. In other words, contrary to popular concerns about echo chambers, social media can actually increase exposure to the outgroup online relative to users’ offline environments. Second, we found that in both offline and online environments, segregation is higher for Democrats than Republicans, meaning that in both online and offline environments, Republicans are more frequently exposed to outpartisans than Democrats. Third, this finding holds for almost all socio-demographic subgroups. Lastly, online partisan segregation is positively correlated with offline partisan segregation.

Our core finding, that American voters are more segregated in their offline environments than in their online environments, resonates with prior literature (12; 30). Our approach departs from this literature in three substantive and methodological aspects: i) it focuses on echo chambers on social media, rather than news consumption, and in a platform heavily used for politics in the U.S. in the time period of our data collection; ii) we use large-scale observational data; iii) it uniquely maps for the same individual their corresponding online and offline partisan segregation. Specifically, whereas the seminal work by (12) focuses on news consumption for an online measure of segregation, and uses only aggregated measures to examine news consumption segregation, we measure online segregation in a way meant to be more comparable to our offline measure: by examining each individual’s online network, not just their media consumption. Compared to (30), we provide evidence of higher offline segregation using individual-level data rather than aggregated county-level measures. We thus avoid the problem of heterogeneity in unit-size that affects county-based analyses. Substantively, we first show that the prior studies’ finding of higher offline segregation extends to the social media space, not only to

news consumption exposure. Methodologically, our study advances this literature with a large-scale, evidence-based approach that directly maps the offline and online social media partisan environments of the same one million voters. Furthermore, our evidence shows that this finding is robust to our large-scale, more granular, contemporaneous individual-level measures — and not an artifact of known biases in survey or aggregate measures of partisan segregation. Lastly, we also present novel evidence, not possible without our measurement approach, that, for the same voter, offline and online segregation tend to move together, with more isolated voters offline also being more isolated online. This pattern is consistent with a recent case study that investigated the geographical embedding of online conversations during the UK’s Brexit campaign (31).

Our findings challenge a common claim in public debate about social media and politics (4; 8; 9; 1; 2; 3), that online environments make citizens more isolated within homogeneous communities. Yet, these claims are often made without using voters’ offline experience as a baseline of comparison. By linking the online networks of Twitter users to the offline residential environments of 1 million U.S. voters, we find that although social media users do live in politically homogeneous networks, these levels of segregation are lower than those observed in offline environments.

Results

Measuring Offline and Online Partisan Segregation

Our results rely on two primary data sources: public voter registration files for over 180 million US voters – which we refer to in the article as *offline data*; and Twitter profiles and their followers’ networks collected through Twitter APIs – which we name in the article as *online data*. The voter file data contains voter registration information up

to 2021, and the online network data collection occurred from January to December 2022. To understand the dynamics between online and offline segregation, we use a matching procedure to match voters with their online profiles on Twitter (32). Using both data sources, we filter users with unique first, last name within the same city-state, and match them accordingly. This results in nearly one million Twitter accounts ² linked to the L2 voter file, which consists of the main dataset we use in the analysis. This matching procedure has been widely used in prior literature (32; 33; 34; 35). For each matched voter, we calculate two segregation measures. ³

Offline Partisan Isolation: For every matched voter, we collect their partisanship ⁴ from the voter file, demographic information, and their residential address. Then, we measure their offline partisan segregation as the proportion of voters’ 1000 nearest neighbors that share the same partisan identity as the matched voter. We call this measure *offline isolation*. When isolation from co-partisan voters is low, we consider voters to live in non-sorted offline environments with high exposure to outgroup voters; when in-group isolation is higher, partisans live in highly homogeneous communities, surrounded by their co-partisans. This indicates that a Democrat (Republican) with high *offline isolation* lives in a more partisan segregated spatial area and is hardly likely to be exposed to a Republican (Democrat) in their offline geographical environment ⁵

²Table 1 in Materials and Methods provides descriptive statistics of demographic information of our final sample, benchmarking it with two other groups: a random sample of Twitter users and a general U.S voting population

³See Materials and Methods for extensive details of these measures, including their mathematical presentation, robustness checks, and details on the data collection

⁴Partisanship is recorded at the time voters register in 30 states and in Washington, D.C. For other states, we use imputed data estimated using procedures developed by (36). We discuss in detail the imputation process in the supplemental materials (SM Section 5)

⁵Similar measures of spatial segregation using individual-level data are often recommended in the specialized literature. These measures are less biased than commonly employed aggregated measures of spatial segregation, which suffer from the modifiable areal unit problem. See (36; 37; 38)

Online Partisan Isolation: Again, for every matched voter (linked to their Twitter profile), we collect all the other Twitter users that they follow (i.e., their friends in Twitter’s terminology).⁶ To estimate ideology for these accounts in the matched user network, we use the ideal-point scores obtained using the procedure outlined in (29). We convert these scores into a partisan group (Democrat, Independent, and Republican) using cutoffs validated with the linked panel for matched voters.⁷ As such, for each matched voter, we have their partisan identity (recorded in the voter file) and the ideology score for each user the matched voter follows on Twitter. Using this dataset, we construct an equivalent measure of online segregation, as the proportion of the matched user’s Twitter friends who share the same partisan leaning as the matched voter.⁸ We call this measure online isolation.

Twitter users often engage with individuals outside of their friends networks. Content may be surfaced through Twitter’s home timeline algorithm, and posts from their friends may not appear in their feed, making the friends network an imperfect estimation of online isolation. To address this concern, we construct the same isolation measure using retweets, quote tweets, mentions, and replies. We find that regardless of the kind of network used to estimate online isolation, our results hold. We provide more details regarding this robustness check in the supplementary materials.

To summarize, both measures represent the relative amount of co-partisans in the

⁶We use Twitter’s terminology for the sake of differentiating between friends and followers. Here, we are assessing “friends,” the set of accounts our matched voters *choose* to follow. However, the use of the term “friend” is not meant to indicate the strength of the tie.

⁷For every matched voter, we have an offline partisanship from their voter file (augmented by the imputation process described in SM Section 5) and an ideology score. Therefore, we can identify valid cutoffs for the ideology scores to assign each partisan group to the users in the matched voter network. We find that 80% of Democrats have scores less than -0.35, and 80% of Republicans in the linked panel are to the right of 0.04. Thus, we classify the friends of the linked panel users as Democrats if their estimated ideology is -0.35 or less and the user as Republican if their estimated ideology is 0.04 or greater. Users in between are considered non-partisans.

⁸We limit the analysis to matched users for whom we can calculate the ideology for at least 10 of their friends. In the supplemental materials (Section 8), we show that this decision does not affect the sample’s median level of online isolation. The main effect of using this cut-off is to avoid users whose networks are not informative.

matched user’s offline and online networks, allowing us to uniquely understand, for the same voter, the relationship between their offline and online segregation. Next, we show the levels of offline and online isolation for the population of matched Twitter users, and the difference between Democrats and Republicans. Lastly, we analyze the differences in online and offline segregation across a wide set of socio-demographic subgroups.

Population-level offline and online isolation

We first describe the distribution of online and offline isolation by partisanship. In Figure 1, Plot A shows distributions of offline (top) and online (bottom) ingroup isolation for Democrats (blue) and Republicans (red). We plot the median (dashed line) for both Republicans and Democrats. Plot B shows the values at the 1st, 10th, 25th, 75th, 90th, and 99th percentiles. These quantities allow us to understand at distinct points of the distribution the differences between online and offline environment for Democrats and Republicans. Within each sub-plot, the top facet corresponds to offline isolation, the bottom row corresponds to online isolation, red represents Republicans, and blue represents Democrats.

We find that Democrats have higher average levels of segregation (more isolated within ingroup communities) than Republicans, both online and offline. And we find that for both groups, the median levels of isolation are greater for the offline distributions than the online distributions. For Democrats, we find that the median offline partisan isolation is 0.71, and the median online partisan isolation is 0.66. For Republicans, we find a median isolation of 0.57 for offline isolation and 0.51 for online isolation.


These numbers also show that Democrats are more isolated offline than are Republicans (note the relative median values of 0.71 versus 0.57), consistent with previous research (36). And these differences also exist in their online environment, where the relative median values are 0.66 for Democrats compared to 0.51 for Republicans. This

indicates that the median Republican Twitter user has in online environment almost equally divided between in-partisans and out-partisans (0.51 in ingroup isolation), which is quite different than the traditional notion of an echo chamber. In SM Figure 11, we also show that differences between Democrats and Republicans are not restricted by the gaps in the median values of online isolation. Republicans also show higher within-group variation, particularly among voters with a large network of friends on Twitter. In other words, among Republicans with high node degree, we detect a wider dispersion, with many voters being overly isolated (embedded in online echo-chambers), while others, high node degree users, are overly exposed to outgroup democrats. Among Democrats, within-group variation is small, and, contrary to the Republicans, as network size increases, Democrats also become more segregated.

Beyond median differences in segregation, Plot B of Figure 1 highlights important dynamics at the extremes of the distributions. At the 90th and 99th percentiles, online isolation approaches offline levels, and partisan gaps narrow. Both Republicans and Democrats show that nearly 90% of their online friends are co-partisans at these extremes. Notably, at the 90th percentile, Republicans experience higher isolation online than offline. These cases illustrate that, at the extremes, online environments *can* resemble echo chambers, with over 80% of one’s network composed of ingroup members. We note that the nature of the online environment is a choice by the individual, and we might expect that some number of people would choose homogeneous environments.

At the lower end of the distributions, however, patterns diverge. Republicans tend to experience lower online isolation than Democrats, with greater exposure to outgroups. By the 25th percentile, Democrats already show majority-ingroup networks (isolation > 0.5), while Republicans at the same percentile remain exposed to liberal users, with about two-thirds of their online ties connecting to Democrats.

A



B



Figure 1 Partisan Offline and Online Isolation. Democrats are shown in blue, and Republicans are in red. Plot A presents the full distribution of values for isolation to ingroup split between Democrats and Republicans. Plot B shows the same distributions across percentiles for one million linked voters.

Correlates between Online and Offline Isolation

We now examine the direct relationship between offline partisan isolation and online partisan isolation. In Figure 2, we show the relationship between offline (x-axis) and online isolation (y-axis). We bin offline isolation values at equal intervals from the variable distribution, compute the average for offline isolation for each bin, and plot the resulting values. Lines on the figure represent a smoothed non-linear function over the averaged points. On the left panel, we plot offline (x-axis) and online isolation (y-axis) for all users. On the right panel, we plot similar information but for Republicans (red) and Democrats (blue) separately.

For all users, regardless of partisanship, we find that as offline isolation increases, changes in the online isolation to ingroup social media users follow accordingly. That is, users with lower offline partisan isolations tend to have lower levels of online partisan isolation. In other words, users with more politically diverse real world neighbors have a more diverse set of online friends (i.e., users followed on Twitter), with a higher presence of outgroup users in their Twitter networks. And we note that this is not an artifact of cross-partisan differences. When we examine Democrats and Republicans separately (panel (b)), we see the same phenomena (though the relationship is not as strong). This finding aligns with the results of (31), which shows that online echo chambers formed in the context of the Brexit campaign in the UK reflect the geographical proximity between Twitter users.

Differences of Online and Offline Isolation Across Voter’s Demographics

We further explore differences in online and offline partisan segregation across demographic and political subgroups using the rich set of variables collected in the L2 voter files. Because offline segregation varies by race, party, population density, and other de-




Figure 2 Relationship between Online and Offline Isolation. Democrats are shown in blue, and Republicans are in red. Each dot represents average online isolation over a hundred binned values of offline Isolation. Lines represent a smoothed non-linear function over the points

mographics (36) , these analyses allow us to see the degree to which online sorting follows offline patterns, and also how the population-wide differences depicted earlier hold for specific socio-demographic groups.

We show how these differences vary across distinct demographic groups. In Figure 3, we show median paired distances between individuals’ offline and online isolation across various sociodemographics and partisanship. Positive values indicate users experience higher levels of offline segregation than online segregation. Across almost all subgroups, including gender, race, age, partisanship, and state type (Red states, Blue states, and Purple (Swing) states) ⁹ users experience considerably higher levels of isolation offline than online. This supports our overall argument that local echo chambers are more pronounced than online experiences are robust across most subgroups. The main exception to this pattern is Republican voters who are older than 60 years. For this sociodemographic group, their levels of online segregation are higher than their offline isolation.

⁹Purple states are defined as states that have a vote-share split between Biden and Trump at 3% or less, meaning that if in 2020, a state’s vote share for Biden compared to its vote share for Trump was within three percentage points of each other, we define that state as a swing state. Using this definition, Arizona, Florida, Georgia, Michigan, Nevada, North Carolina, Pennsylvania, and Wisconsin are classified as swing states. All other states, where the vote share difference is larger than three percentage points either to Trump or Biden, are classified as Red or Blue, respectively.

This is the only group in which their online experience can actually be identified as living in more of an echo chamber online than offline. This finding is interesting considering the consistent evidence that older Republicans are often higher consumers and spreaders of low-quality information on social media (39; 33; 40; 41).




Figure 3 Partisan Offline and Online Isolation Across Socio-Demographic Subgroups. Point estimates represent the median value of the difference between online and offline ingroup isolation for each group with 99% confidence intervals calculated via bootstrapping. Positive values indicate voters who have higher offline isolation than online isolation. Democrats are shown in blue, and Republicans are in red.

In the SM Section 2, we present additional analysis exploring differences in online and offline partisan segregation across distinct levels of urbanization and population density. Differences between online and offline isolation hold among most areas in the United States, regardless of the population density and urbanization. The exception resides in the extremes: in high-density areas, Republicans have more ingroups in their online than offline networks, and in very low-density areas, Democrats show higher online isolation.

In summary, in the extremes of high- and low-population-density areas, social media enables both Democrats and Republicans to connect with other ingroup voters, who are not as easily accessible within their local bubbles.

Furthermore, as discussed in the methods section, our empirical analysis involves several measurement choices that may affect our results. These choices involve the values for the nearest neighbors calculation, the use of ideology scores based on Twitter users' followers' networks, and the cut-off values to classify users as Democrats and Republicans. To validate our results, we provide several sensitivity analyses showing the robustness of the main findings drawn from our study across distinct measurement choices. In the supplemental materials, we demonstrate that the results remain unchanged when we consider only the 500 nearest neighbors when calculating measures of partisan offline isolation. In addition, as in (36), we validate the offline isolation measure when aggregated at geographic units, such as census tracts, precincts, school districts, counties, and congressional districts. We discuss the results at SM, Section 9. We also present in the supplemental materials results using a discrete assignment of partisanship, where the partisan identity is assigned based on the highest probability from the imputation calculation, instead of calculating offline isolation based on the posterior probability for each partisan group. The results remain substantively unchanged, with Republicans being slightly more isolated using this alternative measure.

Lastly, a potential concern related to our proposed measure for *offline isolation* is that an individual's neighbors may not reflect the individuals voters are typically exposed to at work, school, or among their acquaintances, and, therefore, this measure may not accurately reflect the ideological composition of their offline *networks*. To assess this potential concern, we carry out an additional validation analysis that relies on survey data with a representative sample of American voters through YouGov, asking respondents to estimate the number of their neighbors, acquaintances/friends, and coworkers that are

Democrats or Republicans. We then compare measures of segregation for an individuals' networks of their neighbors, acquaintances, and coworkers ¹⁰. In SM Figure 13, we show the relationship between offline isolation measured via neighbor networks (x-axis), and offline isolation measured via acquaintance networks (left plot) and measured via coworker networks (right plot). We find that ingroup isolation measured via network composition of neighbors strongly correlates to network composition measured by both acquaintances and coworkers, indicating that our measure of partisan segregation using the voter file is a reasonable proxy for offline isolation across an extended set of offline interactions.

Discussion

Social media is often seen as a primary driver of the formation of overly homogenous online communities - so-called echo chambers - in which citizens more easily isolate themselves from individuals and opinions from opposing parties and, consequently, reinforce their system of beliefs (4; 42; 5). While a larger body of research has provided a more nuanced view of the prevalence of the phenomenon (10; 11; 13), these studies do not account for partisan segregation in the offline environment. Using a novel dataset connecting online and offline information for close to a million Twitter users in 2022, we deal directly with this concern by explaining how individuals' online and offline environments compare.

Our results make contributions to three distinct areas of scholarly inquiry. First, our findings speak to studies on online echo chambers and social media homophily (29; 10; 42; 26). Prior research in political communication has debated the extent of ideological "bubbles" in social networks, with evidence that partisan homophily drives tie formation online. We contribute to this literature with a direct comparison of online

¹⁰Additional survey details are available in the supporting information (SM Section 1).

and offline sorting. While we do find that social media users curate politically homogeneous communities (consistent with earlier studies), we are able to add the important observation that the degree of ideological isolation online is *lower* than in real-world geographical spaces. Only in the upper part of the distributions do voters experience more of an echo-chamber online than offline. For the most part, social media exposes users to more out-group users than do their offline interactions.¹¹

Second, our results contribute to the emerging literature on integrating offline context with online political behavior, providing novel evidence about partisan sorting both online and offline (12; 29; 10; 42; 30; 31). With rare exceptions (12; 30), most studies rely solely on online information to describe the prevalence of online echo chambers and the role of homophily on online political behavior. To our knowledge, ours is the first study to jointly map the offline and online partisan environments for the same voters at a large scale, bringing the dynamic of offline segregation to this discussion, and engaging with a robust scholarship on spatial partisan sorting in American politics (36; 43; 44; 45). Our approach advances the field by conducting a large-scale comparison of online and offline partisan sorting among a large and diverse set of US voters, showing these patterns tend to move together, and therefore, providing a novel baseline for future research to empirically assess levels of online sorting in the American electorate.

Lastly, our article also speaks to studies focusing on the centrality of partisan geographical segregation in the United States (45; 46; 43). In a recent article, mobilizing the same data and geographic measurement described in our paper, Brown & Enos (36) present evidence of extensive partisan segregation in the US – their results indicate that a large proportion of the American electorate lives with virtually no exposure to out-group voters in their residential environment. This dynamic has been shown to affect elite polarization (47), cooperation for shared benefit across groups (48), levels of trust

¹¹We note that our offline network measure does not directly measure interactions, but our survey analysis (SM Figure 13) suggests it is capturing this.

in government and anti-system attitudes (49), and even adherence to scientific health recommendations during the COVID-19 pandemic (50). By linking millions of voters’ residential contexts with their online networks, our study provides the first behavioral benchmark linking offline geographic segregation to its online dimension, demonstrates that partisan segregation is considerably greater offline than online, and that segregation in both arenas is positively correlated. These findings reinforce how offline partisan sorting is still a decisive feature of American politics even in the years of social media, and provide directions for future research into how these extensive levels of partisan segregation shape online behavior.

Our study is, of course, not without limitations. First, we assume that the measure of offline segregation is predictive of the extent to which an individual is actually exposed to the ingroup and outgroup offline. While we do validate this measure using a survey measure, this measure may miss important avenues of exposure to partisan information. For example, this measure does not include interactions in the offline environment, such as school board meetings, yard signs, community projects, social and religious groups, among others, that could lead to exposure to other partisans or partisan information. Second, our online measure of isolation uses a static version of users’ networks, mostly rebuilding the content users would observe under chronological ranking. Even though part of what users see on their timeline comes from people in their network, at least since 2014, Twitter has provided users with an algorithmic timeline that, among other things, injects non-network content, increases the prevalence of political content (51), recommends more conservative content to users, and facilitates friendship ties with conservative accounts (22). Data limitations in both the offline and online world do not allow us to meaningfully incorporate these algorithmic effects in our measurement strategy.¹²

¹²No data on what users are actually exposed to on Twitter and any other social media platform is available for academic research. This is not a limitation of our work but of any work using social media data available through the Platforms’ APIs. See (41) for an example of *de-facto* exposure data accessed through a partnership between academics and Meta)

With these caveats, we assume that the online and offline measures we employ are roughly equivalent. That is, if an individual is measured to have a 0.33 level of ingroup isolation offline and 0.33 online, then these would lead to proportionally the same amount of exposure to in-partisans and out-partisans. Third, our approach to matching voters to their Twitter profiles oversamples demographic groups that are more likely to have non-duplicated names. We discuss this section issue at length in *Material and Methods*. Importantly, our results show that these demographic groups that are more likely to have duplicated names exhibit higher offline isolation, meaning that, if anything, the main results in the paper are a conservative estimate of offline partisan segregation among social media users in the United States. Despite these limitations, we believe the proxies for the measures that we employ provide valuable insights into the dynamics of offline and online partisan segregation, and this is especially the case when making comparisons within offline and online isolation analyses. We encourage future research to investigate how algorithms might change the online isolation measure we develop, and how users' offline experiences interact with these changes.

Lastly, our study provides a snapshot of the differences between online and offline partisan segregation using a specific platform, Twitter, at a particular moment in time. This consideration is particularly relevant given the significant changes Twitter has undergone in recent years following Elon Musk's acquisition (52; 53). Yet we do not believe this caveat diminishes in any way the importance of our findings. First, our results are temporally valid to a point in time when Twitter functioned as a critical medium for the American public to learn and discuss politics, and as a platform on which politicians and the media relied heavily to communicate and learn about public preferences (54; 51; 55). Second, the vast majority of scholarly work on social media and politics produced in the last decade relies on Twitter data (56), including studies that make claims about the role of social media in facilitating the formation of echo chambers. Therefore, while it is true

that Twitter is no longer the same as it was in 2022, our findings still provide a proper understanding of the social media environment, using a platform that was once critical for political communication.

Our findings complicate the dominant narrative of social media as the primary determinant of whether people end up in “echo chambers.” While online spaces are often portrayed as uniquely isolating, our results show that partisan segregation is actually higher in Americans’ offline environments than in their online ones. For most users, social media provides more, not less, exposure to outpartisan perspectives relative to their local communities. At the same time, we uncover meaningful asymmetries: Democrats remain more isolated than Republicans across both settings, and older Republicans are the only group whose online networks are less diverse than their offline ones. Importantly, the positive correlation between online and offline isolation underscores that the digital and physical worlds are not independent spheres but intertwined environments, with online experiences often reflecting offline realities.

Taken together, these results suggest that the effects of social media on political attitudes and behaviors cannot be understood in isolation from the broader social and geographic contexts in which people live. Rather than treating social media as a standalone driver of polarization, our study highlights the importance of situating digital technologies within the context of offline partisan segregation. This study provides a more nuanced understanding of partisan segregation in the United States and highlights the need for future research to jointly consider online and offline contexts when assessing the political consequences of social media.

Materials and Methods

In this section, we describe each data source in turn, giving special attention to our measures for partisan identity and how we define offline and online segregation measures. We then discuss the research procedures used to connect both data sources before concluding by addressing the ethics of using the data sources together.

Offline Data

For offline data, we use the L2 voter file containing voter registration and demographic information for 180 million U.S. voters.¹³ From the L2 voter file, we extract a voter’s partisanship, demographic characteristics, and home address, the last of which we use to estimate partisan segregation. For each voter, we use the geographic location of their residential address. We measure the spatial distance between a voter and every other voter in the voter’s state using haversine distance.¹⁴ Then, using the K-nearest neighbors algorithm, we identify a voter’s thousand nearest neighbors and consider their 1,000 nearest as voter i ’s offline network. Using these networks, we construct an individual-level measure of offline segregation, estimating the proportion of their nearest neighbors who are from the same partisan group as the matched voters, as in (36). To account for missing partisanship in the voter file, we use Bayesian process imputation, also as in (36). We provide a full description of the imputation process, auxiliary data, and accuracy metrics in the supplemental materials (SM Section 5).

More specifically, using the voter’s residential location, their thousand nearest neighbors, and their partisanship, we measure an individual’s partisan segregation with the following equation:

¹³L2 is a private firm that aggregates individual state voter files, and appends commercially available data to individual records as well as community level publicly available data.

¹⁴This measure is similar to Euclidean distance, but measures distance along an arc rather than directly from point-to-point to account for the curvature of Earth.

$$\text{Offline Isolation}_i = \frac{\sum_{k=1}^{1000} \mathbb{P}(p_k = p_i)}{1000}$$

where k is a given neighbor, i is the matched voter, p_k is the partisanship of the neighbor, p_i is the partisanship of the matched voter. For Democrats, the ingroup partisans are registered and imputed Democrats; for Republicans registered and imputed Republicans. To incorporate uncertainty from the imputation process, $\mathbb{P}$ indicates the probability that p_k is equal to p_i . When the voter’s partisanship is known directly from the voter file, this probability equals one and is equal to the posterior probability from the imputation when voters are independents or non-partisans. We use this quantity as our measure of offline partisan segregation.

Online Data

For online data, we turn to Twitter. In the last decade, Twitter has become a critical platform for disseminating and consuming political content, with a large number of voters and politicians using it as a prominent source of information (54; 51). This is the case, particularly in the United States. While Twitter’s user base has changed more recently, and its prominence for political information has potentially declined after its recent leadership change (57), our data was collected in 2022, before these changes. We collect data from two sources: (1) Twitter’s Decahose API (application programming interface), a ten percent real-time sample of all tweets, and (2) Twitter’s REST API, an API that allows researchers to query both network data for a given user and posts and interactions for a given user. We outline these sources in more detail in the following sections.

Using Twitter’s Decahose to Link Offline Data with Online Data

To match voters with their online profiles, we use data collected from Twitter’s Decahose API, a 10% sample of all tweets collected between January 1, 2022 and December 31, 2022. These tweets contain users’ metadata, including name and self-reported location, which we use for matching. We adopt the method described in (32) to link individuals in the voter file with individuals on Twitter. For the voter file, we take the set of voters with a unique first name, last name, and their location at the city-state level. For Twitter, we follow the same approach, selecting a set of users with unique first names, last names, and city-state. Twitter’s enterprise APIs, including the Decahose, provide derived location information, standardizing location names. For example, if a user had “Manhattan, NY” in their self-reported location field, the derived location would be New York, New York, USA. We use this standardized location field to match to the voter file. From these two datasets, we merge on first name, last name, and location. Similar to (32), we drop any Twitter users that are linked to more than one voter registration and any voters linked to more than one Twitter account, as there are no clear ways to disambiguate true from false positives at scale in this matching procedure. This results in nearly one million Twitter accounts linked to the L2 voter file.

In Table 1, we show descriptive statistics of our sample of Twitter users compared to a survey of Twitter users conducted by the Pew Research Center’s National Public Opinion Reference Survey (58) and the voting population from the 2020 American National Election Survey (59). Compared to the general population, our sample is considerably younger, more male, and more Democratic-leaning than the general U.S. population. These features of our data partially reflect the composition of Twitter users in the United States, as can be seen in Pew Research Center’s National Public Opinion Reference Survey. When compared to Pew Research Data, we find our Twitter panel to have a higher share of men, be more white, be slightly more Democrat, and be slightly older than

the overall Twitter user base. Overall, our matching procedure slightly accentuates the differences between the U.S. population and the Twitter user base. Additionally, as we discuss further in the next section, the higher share of white participants is driven by the matching procedure using unique first and last names within the same city-state.

Table 1 Comparing Demographics: Twitter-L2 Panel vs. Pew Research NPORS 2021 vs ANES 2020

Variable	Sample Values			Δ Twitter Panel		Z-Scores	
	Twitter Panel	Pew	ANES	Pew	ANES	Pew	ANES
Age (Mean years)	41.85(0.02)	39.05(0.76)	48.37(0.34)	-2.80	6.51	-3.6600	19.1200
Gender (% Female)	0.38(0)	0.47(0.02)	0.52(0.01)	0.09	0.14	3.77	17.59
Ethnicity (% White)	0.73(0)	0.66(0.03)	0.66(0.01)	-0.06	-0.07	-2.55	-9.90
Partisanship (% Democrats)	0.47(0)	0.41(0.02)	0.35(0.01)	-0.06	-0.12	-2.44	-15.71

Augmenting Twitter Data with Friend Networks

Using all the matched accounts linked to the voter file, we queried Twitter’s REST API for the friends (i.e. the people the user follows) of each matched user in the dataset. More specifically, in February 2023, we collected the list of accounts each of the matched users follows. For each profile in the matched user network (matched voters’ friends), we measure their partisan identity using the ideal-points method outlined in (29). Notice that this process is only used for the matched users’ friends (for whom we do not have voter file information). Partisanship for matched users is measured using their voter file information.

(29) uses a homophily-based ideology measure, measuring ideology with a Bayesian Spatial Following model that treats ideology as a latent variable inferred by the list of politicians and political elites that a given user follows. We follow the simplified version of the method in (13), which instead uses correspondence analysis to estimate ideology, approximating the results of the full Bayesian estimation with fewer required computing resources. This model results in a scalar value for each individual, ranging from -3 to +4,

where a more negative score indicates that a user is more liberal, while a more positive score indicates that a user is more conservative. For this computation, we use a set of elites, including members of Congress, prominent members of the executive branch (e.g., the President, the Vice President, and various cabinet members), and prominent political organizations (e.g., the RNC and DNC Twitter accounts). We collected the followers lists for each of these elites' accounts. To train the estimation model, we randomly select fifty thousand Twitter users who follow ten or more of the elites. We fit the correspondence analysis model on the correspondence matrix of the fifty thousand users and the elites. Each cell in the matrix indicates one if the user follows that elite and zero otherwise, specifying three component dimensions for the model output. We use the fitted correspondence analysis to project scores into the same ideological space for the remaining users who follow at least three members of the model's elites. We standardize these scores such that the average is zero, and a score of one is one standard deviation of ideology away from zero.

Finally, we assign partisanship to users based on their ideology scores, as described in footnote 7. Using these classifications, we construct a measure similar to our offline isolation measures for users' partisan isolation in their online networks:

$$\text{Online Isolation}_i = \frac{\sum_{k=1}^n \mathbb{P}(p_k = p_i)}{n_i}$$

where k is a given friend of the matched user i , p_k is the partisanship of the friend estimated as described above, p_i is the partisanship of the matched user as in the voter file, and n is the number of Twitter friends from user i for whom we were able to estimate a ideological score. We use this quantity as our measure of online partisan segregation in the following analyses. We limit the analysis to Twitter-L2 matched users for whom we can calculate the ideology for at least 10 of their friends ¹⁵.

¹⁵In the supplemental materials (Section 7), we show that this decision does not affect the median

To summarize, both our measures serve as proxies for the number of voters or social media users in the matched user’s offline and online networks who are co-partisans. In the following analysis, we assume that these measures exist on the same scale, allowing us to compare the levels of partisan segregation for the same voter both online and offline. For example, if a user’s offline isolation is 0.33 and their online isolation is 0.33, we would assume that these lead to equivalent amounts of exposure to and interaction with out-partisans online and offline.

Name uniqueness in the matching procedure

A critical pitfall of this approach arises from excluding potential matches for voters with duplicate first and last names within the same city. This approach raises concerns, given that name uniqueness may correlate with certain demographic characteristics in ways that could affect the estimated levels of partisan segregation. In this section, we therefore probe for these potential selection problems. We first repeat the same matching procedure described above, but for voters with repeated first and last names within the same city. We refer to this group as duplicated-name matches. There is no way for us to verify which of these duplicate matches actually have a Twitter account; therefore, we include all duplicate matches in the analysis below.

Using this set of duplicated-name matches, we first examine demographic differences between this sample and our unique Twitter-L2-matched voters. Second, we estimate offline isolation for all voters with duplicate names. We use the same procedure reported in the paper to measure offline isolation, using the 1000 nearest neighbors to quantify partisan segregation in this new set of voters. Due to Twitter API restrictions imposed after we started writing this paper, we are unable to capture the online networks of these

levels of online isolation. The main effect of using this cut-off is to avoid users whose networks are not informative.

potentially duplicated matches and replicate our measure of online isolation.

Table 2 compares demographic variables of unique matches and duplicate-name matches. We matched 235,105 unique duplicated names with corresponding Twitter profiles, and a total of 1,292,738 voters with non-unique names. The duplicate-name sample is far less white, and that gap is driven mostly by Hispanic voters, who jump from 12.0% to 39.3%, and Black voters, who roughly double in the duplicate-name sample. This sample also leans older, and more male than the sample used in the paper. These compositional changes, the racial ones primarily, move the duplicate-name sample toward having a higher share of Democrats.

Table 2 Comparing Demographics: Unique Name vs Duplicated-Name Matches

	Matched Voters (L2-Twitter)	Duplicated Matched Voters (L2-Twitter)
N	969,995	1,292,738
% White	72.7	37.8
% Black	9.8	18.9
% Hispanic	12.0	39.3
% Other	5.5	3.9
Mean age	41.9	45.3
% Female	37.8	25.8
% Dem	47.3	55.1

Table 3 presents the percentiles of offline isolation for unique matches and duplicated-name matches. This is the critical test for our analysis, as it assesses whether demographic differences across samples affect our outcome of interest. Results indicate that the duplicated-name matches have considerably higher offline isolation in the overall sample. This effect is driven mostly by Democrats. For Republicans, offline isolation is generally lower among duplicated-name matches. These results align with racial differences in segregation discussed in (36), and are therefore not an artifact of our sample. In conclusion, if anything, our procedure for working with unique names provides a lower bound on offline isolation among social media users, thereby alleviating concerns about

bias in the core findings of our paper.

Table 3 Offline partisan isolation: unique-name-matched sample vs. duplicate-name voters on Twitter.

Voters	Percentile of offline isolation						
	1st	10th	25th	Median	75th	90th	99th
Matched Voters (L2-Twitter)							
All voters	0.153	0.371	0.501	0.643	0.794	0.903	0.985
Democrats	0.227	0.416	0.547	0.709	0.855	0.935	0.989
Republicans	0.098	0.307	0.444	0.571	0.681	0.770	0.866
Duplicated Matched Voters (L2-Twitter)							
All voters	0.073	0.339	0.615	0.839	0.932	0.973	0.995
Democrats	0.271	0.583	0.772	0.893	0.950	0.979	0.995
Republicans	0.033	0.114	0.243	0.457	0.656	0.755	0.869

Ethical Considerations for Data Collection

The linkage of Twitter data with voter file data raises privacy concerns for the users contained in the data, particularly as users (and voters) may not expect their data to be linked in this way. As a first step, we only link individuals in the voter file to publicly available Twitter accounts, meaning private Twitter accounts are not included in the data. Second, given the voter file contains fine-grained demographic and residential information, we take seriously considerations of data security and privacy for this data to protect the personally identifiable information (PII) of the individuals contained in the data. Our data management plan was approved by New York University’s IRB FY2021-4869. PII is stored on secure servers with strict access control only allowing read access to individuals on the project. In this manuscript and other outputs of this project, we aggregate the data such that no particular individual can be re-identified from these aggregated statistics.

References

- [1] Haidt, J. Yes, social media really is undermining democracy. *The Atlantic*. Accessed at <https://www.theatlantic.com/ideas/archive/2022/07/social-media-harm-facebook-meta-response/670975/> (2022).
- [2] El-Bermawy, M. M. Your filter bubble is destroying democracy. *Wired*. Accessed at <https://www.wired.com/2016/11/filter-bubble-destroying-democracy/> (2016).
- [3] Hooton, C. Social media echo chambers gifted donald trump the presidency. *Independent*. Accessed at <https://www.independent.co.uk/voices/donald-trump-president-social-media-echo-chamber-hypernormalisation-adam-curtis-protests-blame-a7409481.html> (2016).
- [4] Sunstein, C. R. # republic. In *# Republic* (Princeton university press, 2018).
- [5] Flaxman, S., Goel, S. & Rao, J. M. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly* **80**, 298–320 (2016). URL <https://doi.org/10.1093/poq/nfw006>. <https://academic.oup.com/poq/article-pdf/80/S1/298/17120810/nfw006.pdf>.
- [6] Bail, C. *Breaking the social media prism: How to make our platforms less polarizing* (Princeton University Press, 2022).
- [7] Settle, J. E. *Frenemies: How social media polarizes America* (Cambridge University Press, 2018).
- [8] Pariser, E. *The filter bubble: How the new personalized web is changing what we read and how we think* (Penguin, 2011).
- [9] Conover, M. D. *et al.* Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media* (2011).

- [10] Eady, G., Nagler, J., Guess, A., Zilinsky, J. & Tucker, J. A. How many people live in political bubbles on social media? evidence from linked survey and twitter data. *Sage Open* **9**, 2158244019832705 (2019).
- [11] Guess, A. M. (almost) everything in moderation: New evidence on americans' online media diets. *American Journal of Political Science* **65**, 1007–1022 (2021).
- [12] Gentzkow, M. & Shapiro, J. M. Ideological segregation online and offline. *The Quarterly Journal of Economics* **126**, 1799–1839 (2011).
- [13] Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A. & Bonneau, R. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science* **26**, 1531–1542 (2015).
- [14] Bail, C. A. *et al.* Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences* **115**, 9216–9221 (2018).
- [15] Allcott, H., Braghieri, L., Eichmeyer, S. & Gentzkow, M. The welfare effects of social media. *American Economic Review* **110**, 629–76 (2020).
- [16] Piccardi, T. *et al.* Reranking partisan animosity in algorithmic social media feeds alters affective polarization. *Science* **390**, eadu5584 (2025).
- [17] Levy, R. Social media, news consumption, and polarization: Evidence from a field experiment. *American economic review* **111**, 831–870 (2021).
- [18] Asimovic, N., Nagler, J., Bonneau, R. & Tucker, J. A. Testing the effects of facebook usage in an ethnically polarized setting. *Proceedings of the National Academy of Sciences* **118** (2021).

- [19] Nyhan, B. *et al.* Like-minded sources on Facebook are prevalent but not polarizing. *Nature* (2023). URL <https://doi.org/10.1038/s41586-023-06297-w>.
- [20] Guess, A. M., ..., N. M., Stroud, N. J. & Tucker, J. A. How do social media feed algorithms affect attitudes and behavior in an election campaign? *Science* **381**, 398–404 (2023). URL <https://www.science.org/doi/abs/10.1126/science.abp9364>.
<https://www.science.org/doi/pdf/10.1126/science.abp9364>.
- [21] Allcott, H. *et al.* The effects of facebook and instagram on the 2020 election: A deactivation experiment. *Proceedings of the National Academy of Sciences* **121**, e2321584121 (2024).
- [22] Gauthier, G., Hodler, R., Widmer, P. & Zhuravskaya, E. The political effects of x’s feed algorithm. *Nature* **652**, 416–423 (2026).
- [23] Ventura, T., Majumdar, R., Nagler, J. & Tucker, J. Misinformation beyond traditional feeds: Evidence from a whatsapp deactivation experiment in brazil. *The Journal of Politics* **0** (2025). URL <https://www.journals.uchicago.edu/doi/abs/10.1086/737172>. <https://www.journals.uchicago.edu/doi/pdf/10.1086/737172>.
- [24] Asimovic, N., Nagler, J. & Tucker, J. A. Replicating the effects of facebook deactivation in an ethnically polarized setting. *Research & Politics* **10**, 20531680231205157 (2023).
- [25] Arceneaux, K., Foucault, M., Giannelos, K., Ladd, J. & Zengin, C. Facebook increases political knowledge, reduces well-being and informational treatments do little to help. *Royal Society Open Science* **11**, 240280 (2024).
- [26] Törnberg, P. How digital media drive affective polarization through partisan sorting. *Proceedings of the National Academy of Sciences* **119**, e2207159119 (2022).

- [27] Rathje, S., Van Bavel, J. J. & Van Der Linden, S. Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences* **118**, e2024292118 (2021).
- [28] Brady, W. J. *et al.* Redesigning algorithms to intervene on social norm misperceptions during a national election. *Nature* 1–15 (2026).
- [29] Barberá, P. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political analysis* **23**, 76–91 (2015).
- [30] Tonin, M., Lepri, B. & Tizzoni, M. Physical partisan proximity outweighs online ties in predicting us voting outcomes. *PNAS nexus* **4**, pgaf308 (2025).
- [31] Bastos, M., Mercea, D. & Baronchelli, A. The geographic embedding of online echo chambers: Evidence from the brexit campaign. *PloS one* **13**, e0206841 (2018).
- [32] Hughes, A. G. *et al.* Using administrative records and survey data to construct samples of tweeters and tweets. *Public Opinion Quarterly* **85**, 323–346 (2021).
- [33] Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B. & Lazer, D. Fake news on twitter during the 2016 us presidential election. *Science* **363**, 374–378 (2019).
- [34] Shamir, A., Oser, J. & Grinberg, N. Who is curating my political feed? characterizing political exposure of registered us voters on twitter. *The International Journal of Press/Politics* **29**, 774–794 (2024).
- [35] Baribi-Bartov, S., Swire-Thompson, B. & Grinberg, N. Supersharers of fake news on twitter. *Science* **384**, 979–982 (2024).
- [36] Brown, J. R. & Enos, R. D. The measurement of partisan sorting for 180 million voters. *Nature Human Behaviour* **5**, 998–1008 (2021).

- [37] Reardon, S. F. & Firebaugh, G. Measures of multigroup segregation. *Sociological methodology* **32**, 33–67 (2002).
- [38] Reardon, S. F. & O’Sullivan, D. Measures of spatial segregation. *Sociological methodology* **34**, 121–162 (2004).
- [39] Guess, A., Nagler, J. & Tucker, J. Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science advances* **5**, eaau4586 (2019).
- [40] Guess, A., Aslett, K., Tucker, J., Bonneau, R. & Nagler, J. Cracking open the news feed: Exploring what us facebook users see and share with large-scale platform data. *Journal of Quantitative Description: Digital Media* **1** (2021).
- [41] González-Bailón, S. *et al.* Asymmetric ideological segregation in exposure to political news on facebook. *Science* **381**, 392–398 (2023).
- [42] Bakshy, E., Messing, S. & Adamic, L. Exposure to ideologically diverse news and opinion on facebook. *Science* **348**, 1130–1132 (2015).
- [43] Enos, R. D. *The space between us: Social geography and politics* (Cambridge University Press, 2017).
- [44] Mason, L. *Uncivil agreement: How politics became our identity* (University of Chicago Press, 2018).
- [45] Bishop, B. *The big sort: Why the clustering of like-minded America is tearing us apart* (Houghton Mifflin Harcourt, 2009).
- [46] Levendusky, M. *The partisan sort: How liberals became Democrats and conservatives became Republicans* (University of Chicago Press, 2009).

- [47] McCarty, N., Rodden, J., Shor, B., Tausanovitch, C. & Warshaw, C. Geography, uncertainty, and polarization. *Political Science Research and Methods* **7**, 775–794 (2019).
- [48] Enos, R. D. & Gidron, N. Intergroup behavioral strategies as contextually determined: Experimental evidence from israel. *The Journal of Politics* **78**, 851–867 (2016).
- [49] Cramer, K. J. *The politics of resentment: Rural consciousness in Wisconsin and the rise of Scott Walker* (University of Chicago Press, 2016).
- [50] Baxter-King, R., Brown, J. R., Enos, R. D., Naeim, A. & Vavreck, L. How local partisan context conditions prosocial behaviors: Mask wearing during covid-19. *Proceedings of the National Academy of Sciences* **119**, e2116311119 (2022).
- [51] Huszár, F. *et al.* Algorithmic amplification of politics on twitter. *Proceedings of the National Academy of Sciences* **119**, e2025334119 (2022).
- [52] Bisbee, J. & Munger, K. The vibes are off: did elon musk push academics off twitter? *PS: Political Science & Politics* **58**, 139–146 (2025).
- [53] Barrie, C. Did the musk takeover boost contentious actors on twitter? *arXiv preprint arXiv:2212.10646* (2022).
- [54] Barberá, P. *et al.* Who leads? who follows? measuring issue attention and agenda setting by legislators and the mass public using social media data. *American Political Science Review* **113**, 883–901 (2019).
- [55] Flamino, J. *et al.* Political polarization of news media and influencers on twitter in the 2016 and 2020 us presidential elections. *Nature Human Behaviour* **7**, 904–916 (2023).

- [56] Tucker, J. A. *et al.* Social media, political polarization, and political disinformation: A review of the scientific literature. *Political Polarization, and Political Disinformation: A Review of the Scientific Literature (March 19, 2018)* (2018).
- [57] Cava, L. L., Aiello, L. M. & Tagarelli, A. Drivers of social influence in the twitter migration to mastodon. *Scientific Reports* **13**, 21626 (2023).
- [58] Pew Reserch Center. National public opinion reference survey (2021). <https://www.pewresearch.org/methods/dataset/2022-national-public-opinion-reference-survey-npors/>.
- [59] American National Election Studies. Anes 2020 time series study full release [dataset and documentation] (2021). URL www.electionstudies.org. Version: July 19, 2021.
- [60] Mullen, L., Blevins, C. & Schmidt, B. Gender: Predict gender from names using historical data. r package version 0.5. 2 (2018).
- [61] Imai, K. & Khanna, K. Improving ecological inference by predicting individual ethnicity from voter registration records. *Political Analysis* **24**, 263–272 (2016).
- [62] Enamorado, T., Fifield, B. & Imai, K. Using a probabilistic model to assist merging of large-scale administrative records. *American Political Science Review* **113**, 353–371 (2019).

Offline partisan segregation exceeds online echo chambers in the United States: evidence from 1 million American voters

Supporting Information Files (SIF)

1 Survey Materials

These data are used to validate our measure of offline partisan segregation, which captures the extent to which an individual’s offline network is more ideologically diverse or segregated. Given that the measure we use in the main manuscript is based on an individual’s nearest *geographic* neighbors, it may not represent who they interact with. More specifically, it could be the case that participants live in a geographically diverse area but only spend time with out-partisans. Thus, we ask the following questions of individuals to validate that their neighbors are a reasonable proxy for measuring their offline networks.

We fielded a survey of 3,500 respondents through the survey firm YouGov in early 2023. Respondents were recruited from adult US population, oversampling for social media users. For each respondent, we asked a battery of questions related to the ideological composition of their networks using a grid of questions. These questions were adapted from (12), which estimates the partisan makeup of a respondent’s network. We include figures showing what the participants saw during the survey. In Figure 1, we show the survey question related to the partisan makeup of an individual’s acquaintances. In Figures 2 and 3 we show the same for the partisan makeup of an individual’s neighbors and the people they are acquainted with through work, respectively.



We are going to ask you some questions about all the people that you are acquainted with, meaning that you know their name and would stop and talk at least for a moment if you ran into the person on the street or in a shopping mall. Please answer the questions as best you can.

Of the people you are acquainted with:

	0	1	2-5	6-10	10-20	More than 20	Don't know
How many do you think are Republicans?	☐	☐	☐	☐	☐	☐	☐
How many do you think are Democrats?	☐	☐	☐	☐	☐	☐	☐
How many do you think are neither Democrats or Republicans?	☐	☐	☐	☐	☐	☐	☐
How many are people whose party preferences you don't know?	☐	☐	☐	☐	☐	☐	☐

Figure 1 Survey question for the ideological breakdown of an individual's acquaintances



Now we will ask a similar battery of questions, but focusing on people you are acquainted with through your neighborhood, meaning that you know their name and would stop and talk at least for a moment if you ran into the person on the street or in a shopping mall. Please answer the questions as best you can.

Of the people you are acquainted with in your neighborhood...

	0	1	2-5	6-10	10-20	More than 20	Don't know
How many do you think are Republicans?	☐	☐	☐	☐	☐	☐	☐
How many do you think are Democrats?	☐	☐	☐	☐	☐	☐	☐
How many do you think are neither Democrats or Republicans?	☐	☐	☐	☐	☐	☐	☐
How many are people whose party preferences you don't know?	☐	☐	☐	☐	☐	☐	☐

Figure 2 Survey question for the ideological breakdown of an individual's neighbors



Now, we are interested in learning about people you are acquainted with through work, meaning that you know their name and would stop and talk at least for a moment if you ran into the person on the street or in a shopping mall. Please answer the questions as best you can.

Of the people you are acquainted with through work...

	0	1	2-5	6-10	10-20	More than 20	Don't know
How many do you think are Republicans?	☐	☐	☐	☐	☐	☐	☐
How many do you think are Democrats?	☐	☐	☐	☐	☐	☐	☐
How many do you think are neither Democrats or Republicans?	☐	☐	☐	☐	☐	☐	☐
How many are people whose party preferences you don't know?	☐	☐	☐	☐	☐	☐	☐

Figure 3 Survey question for the ideological breakdown of an individual's coworkers

2 Comparing Online and Offline Isolation Across Urban and Rural Areas and Population Density

Figure 4 presents differences in online and offline isolation across levels of urbanization using core-based statistical areas (CBSAs) to identify major metropolitan areas, and levels of population density¹⁶. Democrats exhibit their highest levels of offline segregation in major metropolitan areas and tracts with high population density. Democrats who are in highly segregated offline environments also tend to be in more segregated online environments.

As shown for the general voting population (not only our sample of Twitter-matched voters), Democrats exhibit their highest/lowest levels of segregation in major metropolitan areas and tracts with high population density (36). We observe similar findings among our sample of matched voters, and more interestingly, these differences closely align with Democrats' online experiences. Democrats, who are highly segregated offline in major metropolitan areas, are also highly isolated on social media. In other words, their online experiences closely match their levels of segregation offline. For Republicans, a similar match appears, but on the other side of the distribution. For those Republicans living in high population density areas, their levels of offline isolation are particularly small (less than 20% of ingroup among their nearest neighbors), as well as their online segregation (around 25% of ingroup in their Twitter networks).

A distinct pattern emerges when looking at differences between offline and online segregation outside of major metropolitan areas and more rural parts of the United States with low population density. In minor metropolitan and outside of metropolitan areas (mostly rural tracts) with a very low population density, the offline isolation of Republicans becomes larger than the offline isolation of Democrats. More interestingly, the

¹⁶Measure of Population Density were adapted from D. H. Montgomery of CityLab. See measures here: <https://github.com/theatlantic/citylab-data/blob/master/citylab-congress/methodology.md>

correspondence between offline and online segregation changes in these rural areas. In these areas, Democrats show higher isolation levels online than offline – a distinct partner from our population estimates. Social media for democrats living in highly Republican areas allows these users to break their “local bubbles“ and give these voters means to communicate with others more aligned with their partisan identity. This is not the Republican experience. Even for voters living in these rural spaces, levels of isolation for republican voters are still smaller than their offline isolation. Even in these highly segregated areas, Republicans are far from experiencing online echo-chambers that matches their offline lives.




Figure 4 Partisan Offline and Online Isolation Across Metropolitan Areas and Population Density. Democrats are shown in blue, and Republicans are in red. The upper plot presents the densities for isolation to ingroup split between Democrats and Republicans across Major, Minor, and Outside Metropolitan Areas using CBSAs. The bottom plot shows the distribution according to population density in each census tract

3 Comparing Online and Offline Isolation Across Subgroups

In the main paper, we provide median differences and statistical tests to compare online and offline isolation across demographic groups (race, gender, and age). In this section, we present the reader with the full distribution for online and offline isolation across these subgroups.




Figure 5 Online and Offline Isolation by Voter's Gender




Figure 6 Online and Offline Isolation by Voter's Age





Figure 7 Online and Offline Isolation by Voter's Race

4 Robustness Checks: Offline Isolation Under Different Cut-offs

In this section, we replicate Figure 1 Plot B from the main paper using a distinct cut-off to calculate offline isolation. In the main article, we use the first 1,000 closest neighbors to estimate offline partisan segregation. Here, we present results using the 500 closest neighbors. Levels of offline isolation increase as we consider this smaller, and geographically closer, set of neighbors.

Number of Nearest Neighbors = 1000



Number of Nearest Neighbors = 500




Figure 8 Partisan Offline and Online Isolation. Democrats are shown in blue, and Republicans are in red.

5 Imputation for Voters’ Partisan Identity

As discussed in the main manuscript, an important challenge of working with information from voter files, compared to survey data, for example, is dealing with voters who are not registered as Democrats or Republicans.

In this section, we describe our procedure and results for imputing partisan identity values for all voters in our augmented data set, as well as their neighbors. Our imputation technique follows (36) using a similar three-step procedure to impute partisanship for matched voters and their neighbors.

In step one, we use partisanship information provided by the voter file vendor L2 to classify all voters. This classification comes from two distinct processes. On one side, thirty states ask voters to declare their party identification in the voter registration. Additionally, for states where party identification is not requested, L2 develops its own modeling strategy to identify partisan identification for every voter in a non-paid state. We use this information from the data provider as our baseline for party identification of the matched voters and their neighbors.

In step two, we consider voters registered with a third party who have a clear left or right ideological leaning as Democrats (left) or Republicans (right). Here, we closely follow the classification scheme provided in (36). Table 4 provides the classification of the third party’s leaning.

Table 4 Classification Code for Third-Party Leaning

Leaning	Third-Party Party
Democrats	Democratic, Green Libertarian, Constitution, Green, Liberal, Progressive, Working Family Party, Peace And Freedom, Socialist, Socialist Labor, Rainbow, Bread And Roses, Worker’s Party, Women’s Equality Party, Social Democrat, Communist, Independent Democrat
Republican	Republican, Libertarian, Conservative, American Independent, Constitutional, Independent Republican
Unknown	Unknown, Non-Partisan, Registered Independent, Independence, Other, Natural Law, Reform, American, Peoples, Declined To State, Patriot, Consumer, Mountain

In step three, we impute the partisan identification for the remaining non-partisan voters through a Bayesian process (36), as described in the equation below:

$$\Pr(PID_i|X_i) = \frac{\Pr(X_i|PID_i) \times \Pr(PID_c)}{\Pr(X_i|R_i) \times \Pr(R_c) + \Pr(X_i|D_i) \times \Pr(D_c) + \Pr(X_i|I_i) \times \Pr(I_c)}$$

The Bayesian process contains two core parts: the individual level likelihood and precinct voter share prior information. The first component represents the likelihood derived from a set of demographics X_i . In this equation, X_{ij} represents the individual-level probability of a voter being part of a certain demographic group conditional on self-reported partisanship (Democrat, Republican, or Independent). These probabilities are estimated using high-quality and large-scale survey data from the Cooperative Congressional Election Study, and use all unique combinations between age groups (18-34, 35-50, 51-62, > 63), gender, and race.

These individual-level probabilities are derived directly from the survey data and are merged back into the profiles recorded for every remaining non-partisan matched voter and their neighbors in the voter files. However, for a small fraction of L2-matched voters,

information about their gender and race is unavailable. For gender, we use information from the census data comparing first name and age on the frequency of male and female names across years to estimate the gender (60). To impute information about voters' race, we use a similar process relying on voters' last names and census information to estimate their racial group (61).


The prior component of the equation is constructed by calculating the precinct-level vote share of Presidential candidates in the 2020 election. Precinct-level vote shares were constructed from data provided by the nonpartisan Redistricting Data Hub (RDH) project, which contains both results and validated precinct boundaries shapefiles. Because the United States does not have a unified coding system for precincts, merging the voter files with precinct-level data requires a set of pre-processing steps. We use a three-step procedure to perform this task. First, we use exact textual matching between the precincts from the voter files and RDH precinct names. Second, we rely on fuzzy textual matching to perform the same task (62). Finally, for all remaining voters whose precincts we could not identify on RDH using textual matching, we rely on geospatial matching between their residential address and the RDH shapefile. This procedure allows us to identify the precinct for more than 98% of all matched sample and their neighbors.

After performing the steps described above, we estimate the quantities described in the imputation equation and impute partisan identity using the largest probability given by the Bayesian process. To assess the accuracy of our imputation, we augment the imputation to also include matched L2 voters for whom we know their partisanship (as provided in step 1 for Democrats and Republicans). Our process achieves an accuracy score of 73.4%, which is comparable to results reported in (36).

6 Robustness Checks: Discrete Assignment of Partisanship

As discussed in the paper, we use the posterior probabilities from the imputation process to calculate the offline isolation measure. This is a more conservative measure in which we incorporate uncertainty from discrete choice models into the isolation measure. However, rather than weighting neighbors by the probabilities, we can just assign voters to a partisan affiliation by using the highest probability calculated from the imputation process. In this section, we present results using the discrete assignment of the partisan identity. Results remain largely unchanged.

Probabilistic Assignment for Partisanship



Discrete Assignment for Partisanship




Figure 9 Partisan Offline and Online Isolation. Democrats are shown in blue, and Republicans are in red.

7 Robustness Checks: Alternative Cutoffs for Online Isolation

To estimate the partisan composition of online networks, we rely on ideal-points methods using as inputs the followers of friends of each matched voter in our panel. As described in (29), this method depends on having users who follow others for which we can have reliable measures of ideology. In our case, to calculate these scores, we use a set of elites, including members of Congress, prominent members of the executive branch (e.g., the President, the Vice President, and various cabinet members), and prominent political organizations (e.g., the RNC and DNC Twitter accounts). We then expand this list to a pre-curated list of 50 thousand Twitter users that follow these accounts, and estimate the ideology score for these users. To estimate the ideology of the friends of our matched voters, we need these users to follow some of these political elites or the Twitter users we have already estimated the scores for previously. On average, we can calculate ideology using this method for 30% of the friends of our matched voters. As described in the paper, to make inferences about the online network of users based on few online friends, we estimate online isolation measures only for those who follow at least 10 other Twitter users for which we can estimate an ideology score. In this section, we demonstrate the robustness of our results to this decision.

We start with some descriptive information of the online network of the matched voters, considering only their friends, we can calculate measures of ideology.

In Figure 10, we show the distribution of the number of friends with ideology scores across the entire matched panel. The x-axis shows the number of friends per matched user with ideology scores, and the y-axis shows the counts of matched voters at each x-axis value. Because this distribution has long tails, we use a log-log transformation on both axes. We split the distribution by party. Two descriptive patterns stand out in these network structures. First, the distributions are heavy-tailed for both Democrats and

Republicans; the majority of matched voters have a few friends (with ideology scores), while a small minority have networks with thousands of friends. Second, Democrats have larger online networks than Republicans, with the median Democrat having 64 friends compared to 44 for the Republican. Importantly, this gap persists across the entire distribution. This discrepancy is expected given that networks are driven by homophily, as our results confirm, and democrats have a higher share of users on the platform. Furthermore, the median matched voter sits well above the at-least-10 threshold we impose for estimating online isolation measures.




Figure 10 Log-Log Distribution for the Number of Friends with Ideology Scores

In Figure 11, we display the mean online isolation conditional on the number of friends for whom we can estimate ideology. This figure shows some interesting descriptive patterns. First, Figure 11 shows that online isolation levels vary conditional on the number of friends, but only among democrats. As network size increases for Democrats, they

become more segregated online. Moving a Democrat from the 25% percentile to the 75% in the node degree increases their online isolation by 0.28SD (0,054 raw units of online isolation), while the same change moves Republicans by 0.01 SD (0.002 raw units). Second, the figure shows a wider dispersion of online isolation among Republicans with a large online network. This within-group heterogeneity is driven by both group size, as the sample of Republicans is smaller than that of Democrats among matched users, and by structural features of the online segregation networks. Republicans with larger online networks show higher variation in their online isolation, with some being overly exposed to Democrats (low isolation), and others being overly segregated in Republican communities (high isolation), and Democrats become consistently more segregated as their network grows. This network feature strengthens our overall finding that Republicans are less isolated online than Democrats: the difference is driven by overall mean differences (as the gap between distributions shows along the entire x-axis) and by greater within-group variation among Republicans with large online networks on Twitter.

Lastly, Table 5 presents differences in node degree across a larger set of demographics, such as age groups, gender, race, partisanship, metro areas, and state partisan leaning. Degree heterogeneity is substantial conditional on age (younger voters have a much smaller network), between White and Hispanic voters, and across partisanship, which aligns with our discussion in the paragraph above. The remaining subgroup differences are more modest.

Lastly, in Figure 12, we replicate Figure 1 in the paper using two alternative cutoffs to filter the data before estimating levels of online isolation. The at-least threshold removes 13.6% (129,639) voters for the full sample of matched voters. However, the paper’s main findings are robust to this cutoff decision, as shown in Figure 12. The latter plot shows that median values of online isolation do not vary across the cutoffs. Importantly, filtering the data with a high cutoff reduces noise in the distribution from extreme online isolation




Figure 11 Mean Online Isolation Across Users with Distinct Number of Friends with Ideology Scores

Table 5 Online Degree Across Subgroups

	N	Median	IQR	% of Panel
Age				
Under 35	314,319	38	14–101	35.5
35 to 60	450,230	70	24–192	50.8
Over 60	108,744	60	19–186	12.3
Gender				
Male	552,850	51	18–141	62.4
Female	330,483	61	21–178	37.3
Metro Area				
Major Metropolitan Areas	788,149	56	19–159	89.0
Minor Metropolitan Areas	52,972	43	16–112	6.0
Outside of Metropolitan Areas	23,812	42	16–107	2.7
Party				
Democrat	547,308	64	21–184	61.8
Republican	338,627	44	16–114	38.2
Race				
White	587,813	59	21–163	66.3
Hispanic	95,694	30	9–91	10.8
Black	78,314	46	14–137	8.8
Other	44,227	53	17–158	5.0
State Politics				
Blue States	350,020	62	20–185	39.5
Red States	324,378	50	18–132	36.6
Purple States	211,537	52	18–145	23.9

Note:

Matched voters with a measurable online degree ($N = 885,935$). Median is of the number of friends with ideology scores; the IQR is the spread of the underlying distribution. ’

values among users with very few friends, removing extreme cases with online isolation at zero or one by filtering the data with a higher cutoff.




Figure 12 Online Isolation Under Different Cutoffs

8 Validating measures of offline segregation with survey measures of acquaintances




Figure 13 Correlation between in-group isolation in an individual’s neighborhood and in-group isolation amongst their acquaintances (left) and their coworkers (right), separately for Democrats (blue) and Republicans (red).

9 Benchmarking offline segregation at aggregated administrative units

In the main manuscript, we measure offline isolation using the 1,000 nearest neighbors of matched voters, following the measurement strategy proposed in (36). (36) argues that this spatial proximity measure of segregation allows for a more accurate measurement of partisan sorting at a granular geographic level in the United States. The authors extensively validate their proposed nearest-neighbor measure, in particular comparing it with approaches aggregated at other administrative units, such as electoral precincts, and ZIP code areas . They conclude that these aggregated measures of geographic segregation consistently underestimate segregation, reducing the levels of partisan isolation as aggregation moves up and the notion of the neighborhood (or offline network) grows (see SI

Figures 9–19).

In this section, we follow (36) and perform a similar validation strategy. We recalculate offline isolation for all matched voters using six distinct measurement strategies for their administrative area: i) their 1,000 nearest neighbors (as reported in the paper), ii) the electoral precinct level, iii) the census tract level, iv) the school district level, v) the county level, and vi) the congressional district level. For i), we use the same formula described in the manuscript; for all others, we apply the same formula but consider all voters from the voter file living in the corresponding administrative unit, for both Democrats and Republicans. We use the same imputation procedure described in the paper for all voters in these units, and we incorporate the uncertainty from the imputation into the offline isolation calculation. We then merge these measures with the corresponding administrative units of the matched voters.

Table 6 shows the values for all six isolation measures for Democrats and Republicans at the 1st, 10th, 25th, 75th, 90th, and 99th percentiles. Our results are robust to alternative definitions; media isolation is largely unchanged at the precinct, census tract, and school district levels. Only as we move up the levels of aggregation to much larger units does isolation tend to decrease. Similar movements occur in the upper tail of the distribution (the most isolated voters, with isolation values above the median). For less segregated voters (1st percentile), isolation increases as the unit of analysis becomes larger. Compared to online isolation, the paper’s core finding is robust to measures at the Precinct, Census Tract, and School District levels, and only reaches similar levels at the larger County and Congressional District units. To summarize, for every unit that more closely matches a voter’s local offline environment (1,000 nearest neighbors, precincts, and census tracts), offline isolation exceeds online isolation for both parties.

These movements are in line with the validation in (36), which the authors describe as an understatement of segregation levels, since it tends to reduce segregation at the

median and in the upper tail of the distribution, and to increase it for those who are less segregated. Given our goal of measuring voters’ local offline environments, as argued in (36), we believe the granular nearest-neighbor measure represents a more accurate description. The median size of these administrative units, presented in Table 6, gives some empirical foundation for this claim: electoral precincts are the closest to the nearest-neighbor measure, but with almost double the size of the 1,000 neighbors. Congressional districts, for example, have a median size of 500,000 voters. Only at these larger units does the offline measure diverge meaningfully from the local environment we aim to capture.

Table 6 Offline Isolation at Administrative Units.

Measure	Median Size	Percentile of offline isolation						
		1st	10th	25th	Median	75th	90th	99th
Democrats								
Nearest neighbors (1000)	1000	0.23	0.42	0.55	0.71	0.86	0.93	0.99
Precinct	1665	0.23	0.42	0.55	0.71	0.86	0.94	0.99
Census tract	3503	0.24	0.42	0.55	0.71	0.86	0.93	0.99
School district	82246	0.25	0.42	0.54	0.69	0.83	0.90	0.97
County	481333	0.26	0.44	0.55	0.65	0.75	0.83	0.91
Congressional district	519694	0.28	0.40	0.50	0.63	0.76	0.88	0.91
Republicans								
Nearest neighbors (1000)	1000	0.10	0.31	0.44	0.57	0.68	0.76	0.86
Precinct	1665	0.10	0.31	0.45	0.57	0.68	0.77	0.87
Census tract	3503	0.10	0.31	0.44	0.57	0.67	0.76	0.85
School district	82246	0.11	0.28	0.43	0.55	0.66	0.74	0.85
County	481333	0.14	0.30	0.40	0.51	0.62	0.72	0.83
Congressional district	519694	0.12	0.32	0.45	0.54	0.62	0.68	0.78

Note: Cells are percentiles of offline isolation at the administrative units for all matched voters. Median Size indicates the median number of voters in the corresponding unit for all voters.

10 Validating Online Ideological Scores

In the paper, we use two distinct approaches to measure the partisan identity. For matched voters, we use their voter file information augmented by an imputation process,

as described in detail the SM Section 5. For the friends of the matched voters on Twitter, we measure their partisan identity using the ideal-points estimation based on their Twitter network (29). This model provides a scalar value for each individual ideological leaning, based on a set of elites, including members of Congress, prominent members of the executive branch (e.g., the President, the Vice President, and various cabinet members), and prominent political organizations (e.g., the RNC and DNC Twitter accounts) that these users follow. We attribute partisanship to users’ friends based on these ideology scores, as described in footnote 7 in the manuscript. (29) provides extensive validation for this method. The authors demonstrate that this approach has high accuracy for both individuals (based on their party registration) and elites, with a correlation between Twitter-based ideal points and DW-Nominate of 0.941 for members of the House and 0.954 for members of the Senate. They also show that this measure is specific enough to generate ideology positions that correlate highly with intra-party measures of ideology. Lastly, this measure has been validated not only in the United States but also in the U.K., Spain, the Netherlands, Germany, and Italy by comparing the Twitter-based ideology estimates to expert surveys (29). In this section, we provide additional validation using two data sources.

We start with *out-of-sample* validation using a general-population panel of YouGov respondents. We asked respondents for their Twitter handles and used this information to estimate their ideology using the same approach and training data as in our paper. This approach gives us ideological scores for 1,549 survey respondents, along with their self-reported partisan identification. Table 7 presents results that validate our approach on this out-of-sample dataset. Across each row, we vary the cutoff used for the discrete partisanship assignment.

Second, we validate our approach using the matched voters data. For a sample of matched voters (N=262,698), we estimated their ideological scores using their friends’

Table 7 Accuracy of Twitter-based ideology cutoffs against self-reported party identification (YouGov donated-Twitter panel)

Cutoff	Voter File Party ID				All respondents	Partisans only	
	N	Precision	Recall	F1	Accuracy	Accuracy	Coverage
Single split at 0 (0.000 / +0.000)							
Democrat	1,031	0.895	0.928	0.911	0.842	0.922	1.000
Republican	384	0.725	0.906	0.806	0.842	0.922	1.000
Independent	134	—	0.000	—	0.842	—	—
80th pct (paper) (-0.351 / +0.045)							
Democrat	1,031	0.913	0.882	0.897	0.820	0.930	0.955
Republican	384	0.731	0.904	0.808	0.820	0.930	0.955
Independent	134	0.179	0.104	0.132	0.820	—	—
75th pct (-0.527 / +0.393)							
Democrat	1,031	0.920	0.840	0.878	0.793	0.957	0.886
Republican	384	0.800	0.867	0.833	0.793	0.957	0.886
Independent	134	0.156	0.224	0.184	0.793	—	—
70th pct (-0.673 / +0.637)							
Democrat	1,031	0.921	0.790	0.850	0.754	0.964	0.829
Republican	384	0.821	0.826	0.823	0.754	0.964	0.829
Independent	134	0.133	0.276	0.179	0.754	—	—

Note:

Nationally representative YouGov panel of respondents who donated their Twitter data. Ideology is the paper’s own score (identical to the values used throughout the analysis). Party ID is the respondent’s modal answer across the eleven panel waves (1-3 Democrat, 4 Independent, 5-7 Republican). Block headings give the Democratic and Republican cutoffs. Because a single split at zero has no middle band it never predicts Independent.

network¹⁷ We can therefore compare their scores with their voter file partisan registration.

Table 8 presents results, and as before, we vary the cutoff used for the discrete partisanship assignment. Table 8 presents these results.

In summary, the validation efforts suggest three findings. First, as reported in (29), this approach is highly effective in recovering measures of ideology that map well into partisan identification in the US. Either with out-of-sample and in-sample validation, the accuracy for partisans is consistently above 80%. Second, the 80th-percentile cutoffs we use in the

¹⁷As we mention in the paper, this is not the measure we use for partisanship of matched voters. For these voters, we rely on their voter file information augmented with the imputation process.

Table 8 Accuracy of Twitter-based ideology cutoffs against voter-file party registration (matched voters)

Group	Per registered party				All matched voters	Partisans only	
	N	Precision	Recall	F1	Accuracy	Accuracy	Coverage
Single split at 0 (0.000 / +0.000)							
Democrat	153,251	0.761	0.871	0.813	0.691	0.853	1.000
Republican	59,469	0.549	0.807	0.654	0.691	0.853	1.000
Other / Non-Partisan	49,978	—	0.000	0.691	—	—	—
80th pct (paper) (-0.351 / +0.045)							
Democrat	153,251	0.780	0.800	0.790	0.668	0.860	0.930
Republican	59,469	0.557	0.801	0.657	0.668	0.860	0.930
Other / Non-Partisan	49,978	0.257	0.103	0.147	0.668	—	—
75th pct (-0.527 / +0.393)							
Democrat	153,251	0.789	0.750	0.769	0.652	0.892	0.837
Republican	59,469	0.629	0.738	0.679	0.652	0.892	0.837
Other / Non-Partisan	49,978	0.266	0.252	0.259	0.652	—	—
70th pct (-0.673 / +0.637)							
Democrat	153,251	0.796	0.700	0.745	0.631	0.904	0.775
Republican	59,469	0.660	0.702	0.680	0.631	0.904	0.775
Other / Non-Partisan	49,978	0.259	0.336	0.293	0.631	—	—

Note: Matched voters carrying a Twitter-based ideology score, scored against their voter-file party registration. Ideology is the paper’s own score (identical to the values used throughout the analysis). Party is the voter’s registration in the L2 file; ‘Other / Non-Partisan’ pools non-partisan, third-party and unknown registrations and has no counterpart in the survey table. Block headings give the Democratic and Republican cutoffs. Because a single split at zero has no middle band it never leaves a voter unlabelled.

paper represent a reasonable balance point between accuracy and coverage. With this cutoff, the scores recover self-reported party for 93% of the YouGov panel and voter-file registration for 86% of matched voters that are wither Democrats or Republicans, while still assigning a label to 93 to 96% of users. Moving the threshold further brings only marginal gains in accuracy. For example, from the 80th to 70th percentile, accuracy improves by 3% among in table 7, and 4% among partisans in 8. However, these tighter cutoffs reduce coverage of the two partisan labels, leaving a larger set of independents in the middle. A single split at zero goes in the opposite direction, and forces everyone to be in a partisan bin. Third, as expected, the classifier does the worst among independents, who are somewhat clustered around zero, but have a more fluid distribution across the full spectrum of ideological scores.

Engagement-Based Metrics of Online Isolation

Our measurement of isolation based only friends is limited and does not count for interactions, content consumption, or algorithmic exposure for the users. While content consumption and algorithm-based exposure cannot be measured outside of Twitter, interactions can be observed through a user’s Twitter timeline. Prior to the closure of Twitter’s Academic API, we collected the last 3,200 tweets (the maximum allowed by Twitter) of each of the respondents in 2023. From these tweets, we measure interactions with other users through retweets, quote tweets, replies, and mentions, which are available in Twitter’s metadata. On Twitter, retweets are when an individual shares another user’s post with no additional message, quote tweets are when a user retweets a post and adds their own commentary, replies are when a user replies to a tweet which shows up under the original tweet, and mentions are when a user tags another user (e.g., @user123 Hello). We parse all tweets for these interactions.

Then, using the same estimation procedure that we use for the friends relationships, we estimate the partisanship of each Twitter user that our panel engaged with. We calculate isolation as: $\frac{\sum_{k=1}^n w_k (p_k = p_i)}{\sum_{k=1}^n w_k}$, where w_k is the number of times an individual interacted with user i , p_k is the partisanship of the individual in our panel, p_i is the partisanship of the user they engaged with. We calculate isolation in five ways. We construct one measure for each type of interaction: a retweet-based measure, a quote tweet-based measure, a mentions-based measure, and a replies-based measure. Lastly, we aggregate over all types of engagement in a single measure.

While content consumption and algorithmic exposure are not directly measurable for researchers outside of the platform, we argue that these engagement-based metrics do, to some extent, capture content consumption and algorithmic exposure. By definition, a user must have consumed a tweet to retweet, quote tweet, or reply to it. Moreover, most of

the tweets a user engages with were likely viewed through the algorithmic timeline, since that is Twitter’s default (as opposed to a reverse-chronological timeline, which consists of the users an individual follows). Thus, the tweets that we measure engagement-based isolation with are some subset of all tweets a user consumed, and they are a subset of what appears in a user’s algorithmic feed.

The following figure contains a comparison of all six measures (the original measure employed in the paper, and the five engagement-based measures that we describe above). The x-axis represents the percentile, and the y-axis represents the isolation at a given percentile, with a separate line for each measure. [14](#) shows that engagement-based measures are highly correlated with the following network-based measure. At the 50th percentile, the measures are about the same. Our findings that users are more isolated offline than online still hold to these distinct measures. For example, as reported in the main manuscript, the median offline isolation is 0.71 for Democrats, and 0.57 for Republicans. For online isolation using the following network information, the median value is 0.66 for Democrats and 0.51 for Republicans, and, when using, for example, online isolation considering any type of engagement, the median values are 0.66 and 0.49 for each partisan group.

However, at the lower end of the distribution, the following network-based isolation measure is higher than engagement-based measures, meaning that users appear more isolated, and at the upper end of the distribution, the following network-based measure is lower than engagement-based measures, meaning our measure makes users at the upper end of the distribution look less isolated. In other words, the distribution of users’ isolation based on engagement relationship is more dispersed, and produces more users with zero contact with the outgroup, as well as more users that are fully isolated. However, the median users are similar across the engagement metrics and the following network metric.




Figure 14 Comparison of different measures for online isolation. Each line represents a different measure: the friend network employed in this paper (black), retweets (green), replies (pink), quote tweets (orange), mentions (blue), and all engagement (yellow). The x-axis represents the percentile, and the y-axis represents the average level of online isolation at that level.