Parametrizing Brexit: Mapping Twitter Political Space to Parliamentary Constituencies

Marco Bastos (City, University of London)

Dan Mercea (City, University of London)

Accepted for publication in Information, Communication & Society

(pre-publication version: some changes still possible)

Abstract

In this paper, a proof of concept study is performed to validate the use of social media signal to model the ideological coordinates underpinning the Brexit debate. We rely on geographically-enriched Twitter data and a purpose-built, deep learning algorithm to map the political value space of users tweeting the referendum onto Parliamentary Constituencies. We find a significant incidence of nationalist sentiments and economic views expressed on Twitter, which persist throughout the campaign and are only offset in the last days when a globalist upsurge brings the British Twittersphere closer to a divide between nationalist and globalist standpoints. Upon combining demographic variables with the classifier scores, we find that the model explains 41% of the variance in the referendum vote, an indication that not only material inequality, but also ideological readjustments have contributed to the outcome of the referendum. We conclude with a discussion of conceptual and methodological challenges in signal-processing social media data as a source for the measurement of public opinion.

Introduction

The referendum on Britain's membership of the European Union was the flashpoint of more than four decades of Euroskeptic politics contesting the country's membership of the supranational organization (Becker, Fetzer, & Novy, 2016). The vote saw those efforts come to fruition as the British electorate was marginally in favour of leaving the E.U., thus opening a new chapter in the political life of the country (Asthana, Quinn, & Mason, 2016) which then embarked on a long process of defining a different relationship with the E.U. In this paper we seek to probe this epochal transformation in British political life by testing whether social media can offer a reliable signal for identifying political alignments as expressed on Twitter. To that end, we provide a proof-of-concept geo-locational analysis of political expression by the British citizenry on Twitter.

Instead of approaching social media analytics as opinion polls, with disputed levels of reliability (Jungherr, Jürgens, & Schoen, 2012; Tumasjan, Sprenger, Sandner, & Welpe, 2011), we examined Twitter data as legitimate manifestations of public opinion in the early 21st century (Anstead & O'Loughlin, 2015), similarly to scholarship investigating the public discourse in pre-industrial bourgeois society of the 18th century that resorted to, and explored extensively, the circulation of information in discursive arenas such as Britain's coffee houses, France's salons, and *Tischgesellschaften* in Germany (Habermas, 1991). As such, the rationale for this study departs from endeavours seeking to forecast the results of the E.U. referendum using social media data as a predictor of voter turnout and party affiliation (Celli, Stepanov, Poesio, & Riccardi, 2016).

In view of the alleged political realignment among Western electorates, we probed into the proposition that not solely material inequality, but also ideological readjustments have contributed to the political outcome of the U.K. voting to leave the E.U. From this perspective, outrage at material inequality has been compounded by a reactionary cultural backlash that has been leveraged and maximized by populist parties and leaders (Inglehart & Norris, 2016). To test this proposition, we devised a conceptual model and a coding scheme to classify content along four ideological coordinates and subsequently trained a dedicated opinion-mining parametric algorithm. We rely on this classifier to analyse a large set of Twitter data collected during the referendum campaign.

Twitter content was collected from a range of hashtags and keywords, including Leave and Remain campaign terms such as #takecontrol and #strongerin and terms that provided a forum for deliberating the referendum (i.e., “Brexit” and “referendum”). Twitter API was also queried to identify the location of users tweeting the referendum. The data we analyse in the following sections thus includes both ideological and geographic markers. We calculated the ideological leaning of users and subsequently mapped them onto voting constituencies in England, Wales, Scotland, and Northern Ireland. As such, the unity of analysis is not tweets or users, but Parliamentary Constituencies from which we model the prevailing ideological landscape as articulated on Twitter in the run-up to the referendum.

In summary, the deep learning algorithm devised for this study is optimized for identifying ideological affiliation, to pinpoint users’ views along a political value space mapped onto Parliamentary Constituencies, and to determine the fit between political expression on Twitter in the period leading up to the vote and the referendum result. In what follows we introduce the conceptual framework underpinning this analysis by unpacking the latent value space before demonstrating its potential for modelling the ideological coordinates of the Brexit debate.

Previous Work

Scholarship informing this study stems from two bodies of literature. Firstly, recent surveys suggest that the British population perceive social media as an important complement to their vote, but they continue to occupy a lower position in the wider ranking of news sources covering elections (Dutton, Reisdorf, Dubois, & Blank, 2017). The sense that social media are nonetheless reshaping the media landscape with momentous consequences for democratic politics flows from the argument that either through a conscious choice or algorithmic filtering, users are narrowly exposed to information that reinforces their political outlook (Sunstein, 2007). Such selective exposure entrenches ideological polarization and forecloses reasoned deliberation (Dahlgren, 2009). While evidence-based treatments of this topic have revealed that exposure to a plurality of political views is likely on social media (Bakshy, Messing, & Adamic, 2015; Fletcher & Nielsen, 2017), social dissemination of political content remains more likely among ideologically similar sources (Barberá, Jost, Nagler, Tucker, & Bonneau, 2015).

Secondly, our research was informed by suggestions of a geographical and socio-demographic patterning of voting preferences in the referendum reported in the U.K. press (BBC, 2016) and scrutinized by academics (Hanretty, 2017; Rennie Short, 2016). The geography of the vote, it was proposed, reflected a socio-economic imbalance between an affluent metropolitan elite clustered in and around London who voted to remain and parts of England and Wales that were economically worse off and voted to leave; and, secondly, a political cleavage between the seat of the U.K. government at Westminster, an increasingly independent-minded Scotland, and Northern Ireland whose economic prosperity and political stability have turned on the existence of an open border with fellow E.U. member, the Republic of Ireland (Rennie Short, 2016).

This study examines public opinion on Twitter against this backdrop of ongoing shifts in deeply engrained ideological leanings (Kriesi & Frey, 2008), which reportedly came to a head in the course of the Referendum campaign (Inglehart & Norris, 2016). We sought to explore whether political talk on social media can quantifiably mirror this process. Specifically, we sought to examine the relationship between communication on social media and the electoral geography of the Brexit referendum to assess the extent to which users tweeting nationalist and populist content would overlap across geographic enclaves; and conversely, whether such pattern could be observed in relation to users tweeting globalist or economist content. In other words, we probed whether Twitter public stream can be used to identify, measure, and model the political consequences of an alignment between the vote and broader ideological orientations expressed by the British public opinion.

Following this line of inquiry, the political geography of the plebiscite was unpicked at the level of local authority areas (Becker, et al., 2016). By means of a best subset selection machine learning protocol for Ordinary Least Square (OLS) regression, Becker, et al. (2016) identified a collection of factors that correlated with the referendum outcome. While contending that a larger turnout in urban areas could have tipped the vote in the other direction, the authors highlighted that the vote to leave correlated positively with a vote for the Euroskeptic UK Independent Party (UKIP) and the British National Party in the 2014 European Parliament elections. Other important correlates of the vote leave were employment in the manufacturing sector, a comparatively lower hourly pay or a higher unemployment rate, the share of rented council housing in the area, longer waiting times for access to the public health system, and lower levels of employment in the public sector.

Demographically, a vote to leave rather than to stay in the EU correlated with the absence of educational qualifications and being 60 years of age or older (Becker, et al., 2016). Cumulatively and in accordance with the economic insecurity hypothesis, the socioeconomic variables were modelled by Inglehart and Norris (2016) in their analysis of the rise of populism in Europe. This supposition—the economic insecurity hypothesis—pertains to a marked decline in the fortunes of the blue-collar working class faced with contracting real incomes, narrowing access to public services such as health, education, housing, or social welfare in advanced post-industrial economies. Their hardship has been attributed to a political inability to spread the economic benefits of an increasingly integrated global economy (Piketty, 2014).

The authors juxtaposed the prevailing economic insecurity hypothesis to the thesis of a cultural backlash against progressive value change (Inglehart & Norris, 2016). Their hypothesis is that socio-economic hardship and resistance to cultural change are mutually reinforced. The result is a cleavage between, on the one hand, the young and well-educated who embraced progressive post-materialist values foregrounding gender, sexual and racial equality, human rights, environmental protection, secularism, and a greater tolerance of migrants. The other side of the divide is occupied by older, less educated sections of the population who experienced a decline in their material conditions, along with the perception of gradual erosion of values associated with industrial societies and solidarity around socio-economic positions, religion, race, and geographic location. This section of the U.K. population saw the cultural politics of identity recognition as a threat to traditional values. Immigration further compounded the disaffection while the EU embodied a cultural threat posed by other European societies which was felt most acutely among people on the lowest education and income levels, manual workers, and the unemployed (McLaren, 2002).

Socio-economic and cultural realignment of British politics

The scholarship reviewed above foregrounds the thesis that not only material inequality led to the result of the referendum, but also a cultural backlash by older, traditional, and less educated voters. This open value competition has augmented political polarization within parties based around cultural issues and social identities (Inglehart & Norris, 2016), a development that maximizes political cleavages and deepens the wedge separating culturally divisive issues. The Conservative Party, in particular, has embraced this cultural cleavage by incorporating a nationalist rhetoric in response to European integration and immigration, devolution, rising secularism, and receding influence in world politics (Kriesi & Frey, 2008).

Political realignment is therefore a process that has been in train for some time, albeit masked by the U.K.'s majoritarian electoral system which kept in place the alternation in government between the two large parties (Dunleavy & Margetts, 2001). The electorate has broadly been divided along two cognitive dimensions: an economic and a cultural one (Kriesi & Frey, 2008). The former was dominant for more than two decades from the 1970s to the early 1990s. The latter became increasingly prominent in the late 1990s and early 2000s. If on the economy voters were divided between supporting or reforming the welfare state, on the cultural dimension they were split between the espousal of, on the one hand, liberal political values, environmental protection, and support for European integration and, on the other, traditional values, "law and order," and a concern with immigration (Kriesi & Frey, 2008, p. 197).

Kriesi and Frey (2008) also showed that Labour and Conservatives converged on a liberal outlook on the economy up to the 1990s, a point of departure when Labour voters became culturally more liberal while Conservatives embraced traditional conservative values. During this period Labour consolidated its foothold among the highly educated and the middle classes, whereas the Conservative Party attracted the least educated and the working classes through a combination of nationalism and cultural conservativism (Kriesi & Frey, 2008), a process heightened by the emergence of populist parties such as UKIP championing traditional values alongside “nationalistic and xenophobia appeals, rejecting outsiders and upholding old-fashioned gender roles” (Inglehart & Norris, 2016, p. 30).

Coordinates of the Brexit Ideological Value Space

The conceptual model comprises two axes opposing globalism to nationalism and economism to populism. There are three notable definitions of populism that have guided social science research (Mudde & Rovira Kaltwasser, 2012). Firstly, populism amounts to a movement galvanized by a hypnotic leader that crosses class boundaries; it is, secondly, a manner of doing politics that mainlines the relationship between political leaders and the electorate to the detriment of political parties. Thirdly, populism is a political discourse premised on the claim of greater authenticity in the representation of the experiences and beliefs of an oppressed majority, who sits in antagonism with a hegemonic minority (Laclau, 2005). These definitions have been thrown into question as political parties started to adopt a catch-all response to the erosion of their relationship with the electoral base (Kriesi & Frey, 2008), a process that conflated populism with demagoguery but also took it to the center stage of party politics.

Despite these developments that limited the analytical value of the concept to the point of being employed as a floating signifier, the three definitions of populism pivot on the variance between “the people” and “the establishment” (Mudde & Rovira Kaltwasser, 2012). Asserting socio-economic, political, and cultural alienation, populist discourse further expresses this antagonism as geographical distance separating, for instance, loci of political and economic power such as Brussels, London, or the South-East of England from the rest of the U.K. (Wills, 2015). The latter embodied a privileged and rootless cosmopolitan elite (Bauman, 2012)—corporate or governmental—vilified as panderers to economic and political globalization personified by the E.U. and European integration (Kriesi, Grande, Lachat, Dolezal, Bornschier, & Frey, 2008; Woods, 2009).

Nationalist parties have channeled this disenchantment by capitalizing the skepticism of economic trade agreements, technological disruption, and the belief in a culturally and ethnically monolithic state. This homogeneous territory would, firstly, reassert its authority in the face of a perceived abdication of economic self-interest in trade liberalization agreements and, secondly, exert social control over labor migration flows that placed crippling pressure on the institutions of the welfare state (Mudde, 2000, 2004). Conversely, globalism, the third coordinate of the latent ideological space pertains to a rights-based universalistic worldview that regards individual citizens as free agents operating in a global economy of increasingly convergent national political systems (Turner, 2002).

To summarize, in its more basic forms, populist messages advance a discontent directed at elites and the establishment and foreground popular will, while nationalist sentiments revolve around notions of national exceptionalism, sovereignty, and nativism (Parker, 2016). In opposition to that, a prevalent response to popular disenchantment has been a drive towards greater efficiency in economic policy and analysis, policy-making, and government administration. This is the fourth coordinate in the latent ideological space which we term economism, conceived in opposition to populism and that emphasizes consensus building, due process of law, and accountability, but also expert analysis and evidence-based policy making that could drive consensus across ideological fault-lines (Nilsson & Carlsson, 2014). As such, economism refers to the comprehensive political consensus to safeguard free market economics embodied in government policy and the array of expert bodies—from the Bank of England to think-tanks, business, and trade organizations—which have helped define and uphold it in the last three decades (Crouch, 1997). Figure 1 shows the ideological coordinates and the political value space that serves as the analytical baseline for this study.

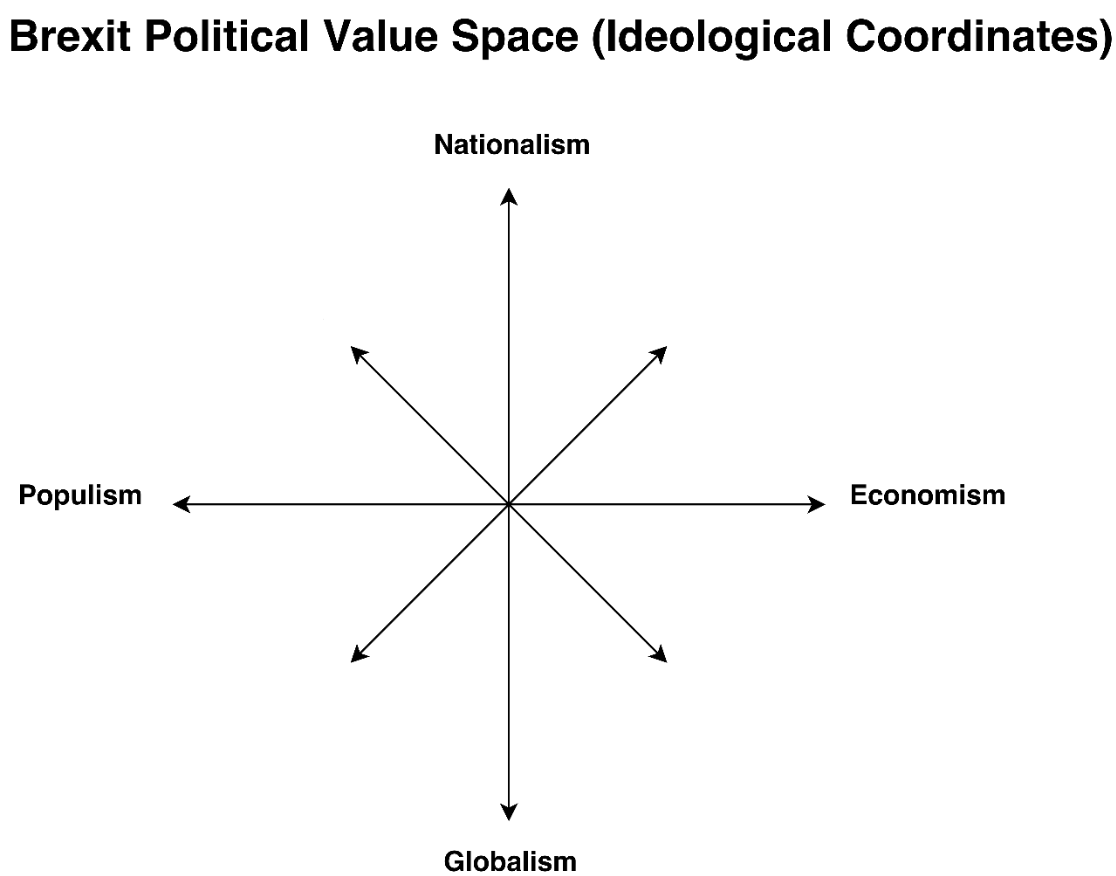

Figure 1: Ideological coordinates of British public opinion and political value space

It is against this backdrop of political realignment that the use of social media and data analytics was portrayed as having facilitated accurate canvassing by, for example, the anti-establishment Vote Leave campaign (Cummings, 2016). Equally, social media data have been used to model the result of the vote on June 23rd based on agreement with the Leave or Remain campaigns expressed on Twitter (Celli, et al., 2016). Accordingly, we sought to model the political debate on Twitter in the weeks leading up to the vote and subsequently mapped the value space onto Parliamentary Constituencies (BBC, 2016; Becker, et al., 2016; Hanretty, 2017; Rennie Short, 2016). We tested the hypothesis that an economist and globalist discourse would cluster around affluent metropolitan areas with a higher-than-average concentration of groups who have reaped the economic, political, and social benefits of globalization. We concurrently test the hypothesis that economically fragile northern Britain would more readily embrace nationalist, anti-immigration, and populist claims which amalgamate protectionist calls for Britain to shield itself from the global economy, to control flows of people, and to regain its national sovereignty from the E.U.

More specifically, we test the hypothesis that the distribution of users advocating either side of the campaign mirrors the results of the referendum (H1). Further, we anticipate the latent ideological space underpinning the classifier maps onto the referendum results across Parliamentary Constituencies, namely along the Globalism-Nationalism polarity (H2), and the Economism-Populism axis (H3). In close relation, we conclude our analysis by modelling the dependent variable RemainPcnt (H4)—the percentage of vote of the Remain campaign—with a multiple regression model that incorporates the abovementioned independent variables (RemainLeave, GlobNat, and EconPop) along with demographic variables retrieved from publicly accessible census data similarly mapped onto Parliamentary Constituencies (ONS, 2011). In sum, the four central hypotheses tested in this study are the following:

H1. The distribution of tweets advocating either side of the campaign matches the vote results across Parliamentary Constituencies (hashtags);

H2. The distribution of Nationalist and Globalist tweets matches the vote results across Parliamentary Constituencies (classifier);

H3. The distribution of Populist and Economist tweets matches the vote results across Parliamentary Constituencies (classifier);

H4. Tweets mapped onto the ideological value space, combined with demographic variables, can account for geographic heterogeneity in the referendum results (multiple regression model).

Data and Methods

For the purposes of this study we relied on the Twitter Streaming and REST APIs to collect a total of 8,821,116 tweets using a set of keywords and hashtags, including relatively neutral tags such as brexit, referendum, inorout, and euref, but also messages that used hashtags clearly aligned with the Leave campaign: voteleave, leaveeu, takecontrol, no2eu, betteroffout, voteout, britainout, beleave, iwantout, and loveeuropeleaveeu; and hashtags clearly aligned with the Remain campaign: strongerin, leadnotleave, votein, voteremain, moreincommon, yes2eu, yestoeu, betteroffin, ukineu, and lovenotleave. Vocal hashtags supporting the campaigns are leveraged to identify messages advocating each side of the referendum: The Vote Leave or Vote Remain campaigns. We subsequently removed messages tweeted before 15 April 2016, the starting date of the official campaign period, and 24 June 2016, the end of the referendum campaign. Messages flagged as likely to have been tweeted by bots were also removed (Bastos & Mercea, 2017).

Next, we queried the Twitter REST API to retrieve the profile of users that tweeted the referendum. We managed to retrieve 95% of the user profiles that appeared in the data (794,949 out of 834,878). Profile information, along with information tweeted by the users, was pivotal to identifying the location of the user base. We triangulated information from geocoded tweets (subsequently reverse-geocoded), locations identified in their user profile (then geocoded), and information that appeared in their tweets. The triangulation prioritizes the signal with higher precision, hence geocoded information is preferred if present. When not available, we look at the location field in users' profiles and geocode that location. If neither source of information is available, we check for information in their tweets, but only in cases where the *place_id* field of the API response returns relevant information.

As a result, a considerable portion of user locations in our dataset could be identified only to city or postcode level. Nonetheless, we succeeded at identifying the geographic location of 60% of users that tweeted the referendum (482,193 out of 794,949) who form our population of interest. From this cohort of 482,193 users tweeting the referendum, only 30,122 were based in the U.K. Upon identifying the location of users, we removed user accounts located outside the United Kingdom or whose location we could not identify up to postcode level. This reduces our dataset to 565,028 messages or 11% of all collected messages; a sample of messages that is sufficiently large to allow for testing the hypotheses underpinning this study.

Campaign Advocacy

For each tweet, we count the number of hashtags advocating the Leave and Remain campaigns. We tag the message as Remainer or Leaver based on the highest number of vocal hashtags used in association with each side of the campaign. Messages without hashtags advocating either side are tagged as Neutral. The frequency count is aggregated and used to calculate the affiliation of users that tweeted or retweeted hashtags advocating either side of the campaign. Highly polarized messages—i.e., tweets including several supporting hashtags—are however uncommon. For users championing the Vote Leave campaign, only 16% of their messages included more than one such hashtags. These messages are yet more uncommon in the vote Remain campaign, where only 2% of messages included more than one hashtag clearly associated with that side of the campaign.

We conclude the identification of users campaigning for either campaign by calculating the mode or "mean campaign affiliation" of users based on the frequency of campaign-supportive hashtags used throughout the period. The mean affiliation of users can only be calculated for accounts that actively participated in the referendum campaign on Twitter. In other words, only users that actively tweeted or retweeted content clearly aligned with one side of the campaign are identified in this step of the data processing. We believe this approach, grounded on the mean affiliation per user, reflects strong campaign membership with low probability of false-positives and below we detail how this measure compares with the variable returned by the machine learning algorithm.

Brexit Classifier

The Brexit Classifier is a machine learning algorithm that resulted from multiple tests to identify the four key ideological coordinates explored in this study. We relied on two expert coders who classified 10,000 tweets along the ideological coordinates of Globalism, Economism, Nationalism, and Populism. We controlled for intercoder reliability by double-coding a random sample of tweets (*N*=100) repeatedly, and after four rounds we achieved a Krippendorff's alpha of 0.94 for the complete value space, with alpha of 1 for the Globalism-Nationalism dyad and 0.86 for the Economism-Populism polarity. We relied on this trained set of tweets to parametrize the machine learning algorithm using text vectorization (Selivanov, 2016), an approach purposefully-built for text analysis.

Unlike frequency-based approaches to text classification, which simply compute the number of positive and negative words (or hashtags) and draw a conclusion based on the final sum, text vectorization is a deep learning algorithm that draws context from phrases. It is often deployed to analyze and classify large text corpora, including user feedback, reviews, and comments. The algorithm can handle linguistic variation and performs well with misspelled or poorly constructed sentences, a marker of Twitter communication, because it considers the entire body text of tweets to infer ideological inclination. It is independent from hashtags, though in the Brexit corpus we found hashtagged tweets to be more vocal and likely to display a clear alignment with one of the four ideological coordinates. As a result, and unsurprisingly, the algorithm consistently identifies campaign hashtags as valid indicators of tweets ideologically leaning towards a given position in the political value space.

Training a machine learning algorithm is fundamentally a trade-off between recall, the number of correct results divided by the number of possible results, and precision, the ratio of positive and relevant matches. In other words, the more variables the algorithm has to identify (in our case there are four: globalism, economism, populism, and nationalism), the higher the likelihood that the algorithm will be unsuccessful. For the purposes of this study, the algorithm needs to identify at least one and a maximum of two ideological coordinates, as the polarities globalism-nationalism and economism-populism are mutually exclusive. Given the disjoint assumption of the political value space, we maximized precision and recall by splitting the ideological value space along two polarities and training two separate algorithms later combined into a single classifier (i.e., the Brexit Classifier). This approach successfully returned substantially more relevant results while also returning most of the relevant results.

We relied on the abovementioned set of 10,000 manually coded tweets to assign a value (positive or negative) to each of the concepts we have sought to map, with the algorithm calculating the probability of positiveness and negativeness for each ideological polarity (Globalism vs. Nationalism and Economism vs. Populism). For each ideological pair, the classifier returns a range of values from 0 (completely globalist) to 1 (completely nationalist), so that values from 0.45 to 0.55 are somewhere in the middle of this scale and assumed to be relatively neutral. The algorithm was trained using Document-Term Matrix (DTM), vocabulary-based vectorization, and the TF-IDF method for text preprocessing. Figure 2 shows the area under the curve on train and test datasets for the Economism-Populism and Globalism-Nationalism ideological pairs (AUC=0.8697 and AUC=0.901, respectively). The algorithm performed well for the set of 565,028 tweets explored in this study and we expect it to perform reasonably well in other national contexts in which nativist and populist sentiments might be emerging. In the last step of the classification, the algorithm calculates the best fit, projects the results along spatial coordinates comprising the four ideological dimensions, and estimate significant oscillations between any of the ideological pairs.

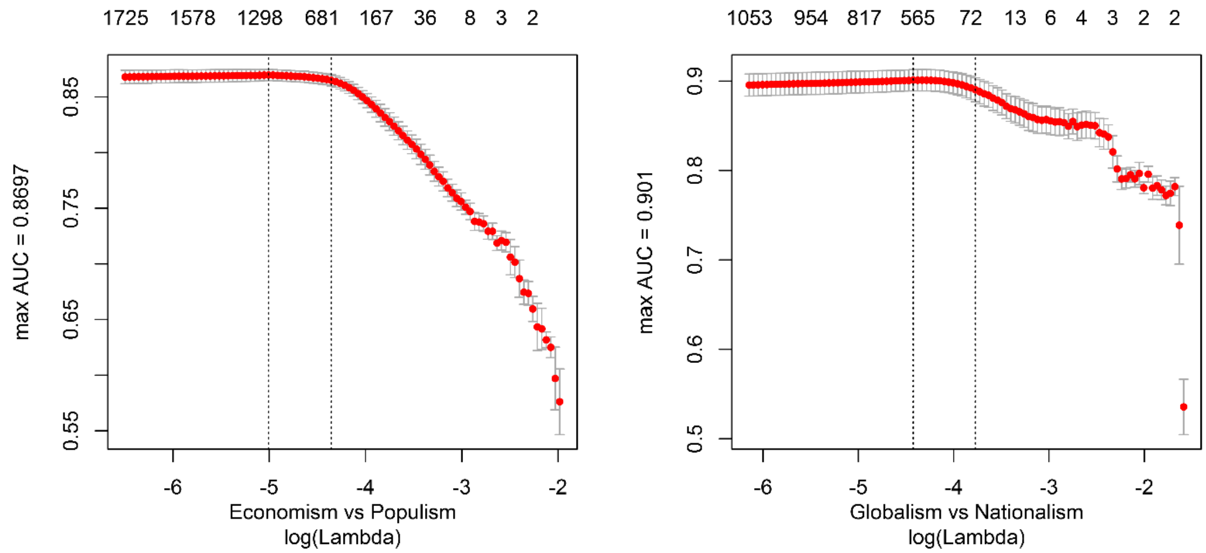

Figure 2: Area Under the Curve calculated from train and test datasets for the Economism-Populism and Globalism-Nationalism ideological pairs

Unit of Analysis: Council Wards and Parliamentary Constituencies

To leverage the granularity of our data, we rely on previous research that successfully mapped the United Kingdom's referendum on membership of the European Union—restricted to local authority level—to parliamentary constituency level using a scaled Poisson regression model that incorporates demographic information from lower level geographies. This approach relies on a principled method of areal interpolation to aggregate the results at ward or constituency level, along with voting estimates at the level of council wards for authorities that have not disclosed the results at such granular levels (Hanretty, 2017; Huyen Do, Thomas-Agnan, & Vanhems, 2015). The processed referendum data is thus relatively granular with data down to the ward level in England, Scotland, and Wales. As the ward system does not exist in Northern Ireland, the data were aggregated at the Local Authority District, thus overcoming inconsistencies between local authorities and successfully mapping postcodes to Parliamentary Constituencies. In short, we adopt ward level data when available and estimates of the referendum results where ward level data were not made available by the authorities. Such estimates advanced by previous research (Hanretty, 2017) allow us to investigate the extent to which the geographic distribution of tweets supporting each side of the campaign and voicing opinions attached to the ideological coordinates mapped in this study interact with the how constituencies voted in the referendum.

Mapping Twitter Data to Council Wards and Parliamentary Constituencies

Mapping geographically-rich social media data onto census area or electoral districts is challenging due to the hierarchical subdivision of U.K. local government areas into various sub-authority areas and lower levels such as enumeration districts. As council wards comprise the most granular level to which we could retrieve results or estimates for the referendum vote, we sought to map referendum-related Twitter activity to this unit of geographic analysis. Therefore, we geocode and reverse-geocode the location of users that tweeted the referendum and subsequently match postcodes to wards and Parliamentary Constituencies using the database provided by National Statistics Postcode Lookup (ONS Geography, 2011). Twitter users are thus simultaneously matched to the fields OSLAUA, OSWARD, and the PCON11CD (Local Authority, Ward, and Constituency codes, respectively). The first field includes Local Authority District (LAD), Unitary Authority (UA), Metropolitan District (MD), London Borough (LB), Council Area (CA), and District Council Area (DCA). Where the council ward system does not exist (i.e., Northern Ireland), data were aggregated using these authorities to cover the entirety of the United Kingdom.

Upon geocoding the self-reported location of users, we found that only 30% of them were based in the U.K., with 19% of users that participated in the Brexit debate based in the U.S. and nearly 30% in other E.U. countries. Also surprising is the large geographic spread of the British Twitter user base, with London accounting for 14%, Lancashire 7%, Kent, Essex, West Yorkshire, and West Midlands ranging 3-4%, and South Yorkshire, Hertfordshire, Cheshire, Merseyside, Surrey, and Hampshire at 2% each. Taken together, each of these geographic groups are of comparable size to London in the share of users that tweeted the referendum.

We ultimately consolidate referendum and Twitter data based on OSLAUA (Local Authorities) and PCON11CD, which is the standardized ID code for each Parliamentary Constituency, the only GSS (Government Statistical Service) beyond European electoral region that is available for Northern Ireland and is consistent across the four countries included in the United Kingdom (ONS Geography, 2017). Using postcode as the common geographic marker across databases, this last step of data aggregation allows for pairing Twitter and referendum data based on Local Authority District, each comprising a range of postcodes. We assigned pseudo codes when no postcodes or grid reference were made available by the authorities, particularly in the cases of the Channel Islands and Isle of Man. Data provided by the Office of National Statistics assigns the range E06 (UA), E07 (LAD), E08 (MD), and E09 (LB) to England; W06 (UA) to Wales; S12 (CA) to Scotland, and N09 (DCA) to Northern Ireland, with the pseudocodes L99 being assigned to Channel Islands and M99 to Isle of Man. Following these procedures, we first calculate the user-average score returned by the Brexit Classifier (Globalism, Economism, Nationalism, and Populism) and the mean campaign affiliation based on advocacy hashtags tweeted by users.

In summary, we calculate the mean campaign affiliation, mean globalist-nationalist, and mean economism-populism for each user that tweeted the referendum. Lastly, we match users to Local Authority Districts to test the hypotheses driving this study. Twitter data is therefore aggregate first at user level, and subsequently at constituency level, which is the unit of analysis employed in this study. The resulting dataset includes multiple streams of Twitter data consolidated into a single database of online and off-line activity at the constituency level: firstly, their ideological expression on Twitter, and secondly, their voting preferences relative to the 2016 U.K. E.U. membership referendum.

Limitations of the Methods and Data

There are important limitations associated with the ideological polarities developed for this study. Firstly, during the process of training the classifier we struggled to separate economism from populism, as many of the populist claims are economic in nature. This is reflected in the lower AUC score for Economism-Populism compared with Globalism-Nationalism. We addressed this challenge by accentuating the policy and expert-oriented component of the economism polarity, which sits in opposition to populist views that appeal to emotion and the perceived rights “of the people,” a value that is difficult to unpack but that stands visibly against the value space occupied by economism. Secondly, the clear identification of messages with nationalistic content has limited heuristic value, as nationalistic sentiments in Scotland and England refer to fundamentally different political agendas. Lastly, the various sampling techniques applied to the data, particularly the geographic rendering of user locations up to postcode level, reduced the universe of collected tweet to 11% of the dataset. We expect this set of data to offer a defensible representation of the Brexit debate and our conclusions are conditional on these constraints.

Results

The dependent variable percentage of vote Remain (henceforth RemainPcnt) is positively correlated with the independent variables tested in this study: Globalism-Nationalism (henceforth GlobNat), Economism-Populism (henceforth EconPop), and RemainLeave, i.e., support for either campaign measured by advocacy-hashtags (r=.31, .46, and .26, respectively). The independent variable GlobNat refers to the polarity Globalism and Nationalism, in which a message with the highest Globalist content is rated 1 and a message with the highest Nationalist content is rated 0. The same scale applies to the variable EconPop, with the polarities Economism and Populism varying between 1 and 0, and RemainLeave, which is the average affiliation calculated per user, also normalized to a scale of 1 to the Remain campaign and 0 to the Leave campaign.

We approach hypothesis H1 by exploring the similitude between the geographic distribution of politically-charged tweets and the geography of the vote, aggregated to the level of Parliamentary Constituencies. As detailed in the Methods section, we relied on hashtags unequivocally advocating either side of the campaign to generate a vector varying from 0 (total Leave support) to 1 (total Remain support). This variable (RemainLeave) is significantly correlated with the results of the vote (r=.26. p<.0001), but upon regressing the referendum results variable (RemainPcnt) on this explanatory variable we found that it can explain a modest 6% of the variance found in the data (R^2_{adj}=.068, p=2.377e-11). Therefore, we reject hypothesis

H1 and conclude that the spread of hashtags advocating either side of the campaign is a poor predictor of the referendum outcome at granular levels such as Parliamentary Constituencies.

Next we tested hypotheses H2 and H3 by investigating whether the distribution of Nationalistic and Globalist tweets mirrored the distribution of Remain and Leave vote across Constituencies. The two explanatory variables are significantly correlated with the results of the referendum (r=0.31 and 0.46, for GlobNat and EconPop, respectively, p<.0001), but it is the results of the regression that are particularly interesting. While the polarity Globalism-Nationalism has a modest albeit significant explanatory power (R^2_{adj}=.10, p=9.226e-16), the polarity Economism-Populism explains over one fifth of the variance found in the results of the referendum (R^2_{adj}=.21, p<2.2e-16). Figure 3 shows how the classifier positioned each of the half a million tweets processed in this study with the fitted line representing the trend detected by the algorithm.

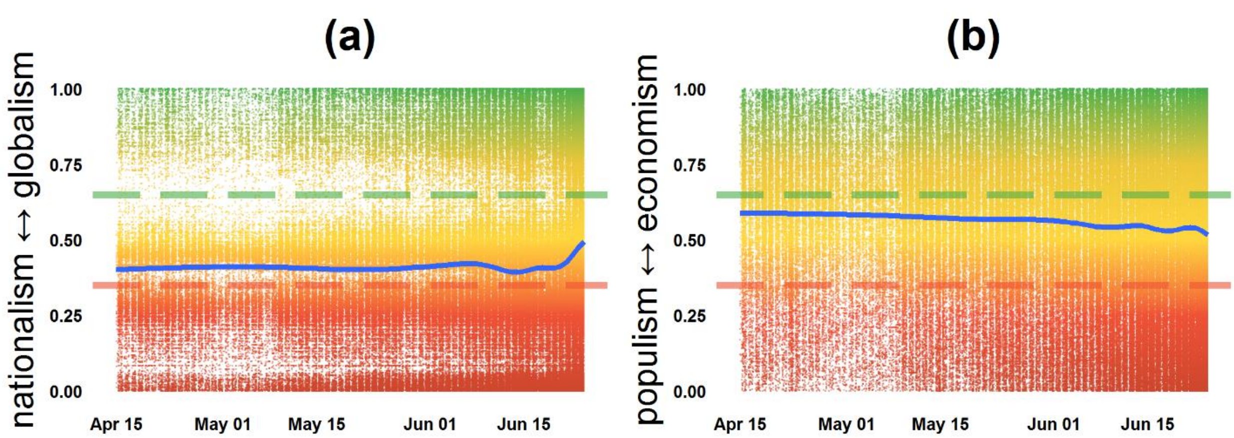

Figure 3: Ideological value space calculated from Twitter messages. Blue line indicates the probability of nationalist versus globalist (a) and populist versus economist (b) sentiments, respectively. Plotted dots indicate the position of each of the half million messages

The classifier reported a strong nationalist sentiment in the data, which persists throughout the campaign and is only offset in the last days when a globalist upsurge brings the British Twittersphere closer to an equally partitioned divide between nationalism and globalism. For most of the campaign, the overall sentiment is decidedly nationalistic averaging .40, which translates to three quarters of messages having a nationalistic sentiment. On the Economism vs Populism spectrum the sentiment is reversed: most messages tweeted in the period (61%) are preoccupied with economic implications of the decision to leave the E.U. Though messages with a strong populist appeal account for less than 40% of the total messages, the trend shown in Figure 3 is of growing occurrence of populist messages in the weeks and days leading up to the vote, with messages centered on economic issues moving out of the debate as a populist discussion balloons.

Lastly, we test hypothesis H4 that the ideological value space can be combined with demographic variables to model public opinion formation relative to the referendum results (multiple regression model). We begin by regressing the vote results on the variables RemainLeave (support with vocal hashtags), the variables GlobNat and EconPop generated by the classifier, and the percentage of the population that is economically active in each Parliamentary Constituency. Except for GlobNat, all variables are deemed significant and the model explains about one quarter of the variance found in the results of the referendum (see Table 1). Next, we include a range of demographic variables and perform a stepwise model selection by Akaike Information Criterion (AIC). The returned stepwise-selected model includes an ANOVA component that rejects the variable GlobNat (due to low significance) and incorporates the variables unemployment, valid votes, electorate, and retired population, which unsurprisingly much improve the model. Electorate size and valid votes are variables that favor urban areas (apart from London) where Twitter penetration is considerably higher (Pew Research Center, 2013), while economic variables such as unemployment and size of the retired population have been found to be associated with vote Leave (Becker, et al., 2016).

Table 1: Models incorporating demographic variables, such as unemployment and economically active population, with the ideological value space on Twitter to explain the referendum results

(a)

Residuals:

Min	1Q	Median	3Q	Max
-0.30132	-0.06142	0.01312	0.07234	0.23328

Coefficients:

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	3.664170	1.162669	3.152	0.00170 **
RemainLeave	-4.255821	2.525347	-1.685	0.09244 .
GlobNat	0.124783	0.327011	0.382	0.70290
EconPop	-1.540038	0.173522	-8.875	< 2e-16 ***
econActivePcnt	-0.003563	0.001174	-3.035	0.00251 **

Signif. codes:
0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
Residual standard error: 0.1004 on 625 degrees of freedom
Multiple R-squared: 0.2317, Adjusted R-squared: 0.2268
F-statistic: 47.12 on 4 and 625 DF, *p*-value: < 2.2e-16

(b)

Residuals:

Min	1Q	Median	3Q	Max
-0.298907	-0.052397	0.008676	0.061566	0.225134

Coefficients:

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	1.915e+00	5.156e-01	3.714	0.000222 ***
RemainLeave	-2.385e+00	1.092e+00	-2.183	0.029404 *
EconPop	-9.355e-01	1.400e-01	-6.681	5.26e-11 ***
unempPcnt	2.734e-02	3.444e-03	7.939	9.50e-15 ***
Valid_Votes	1.975e-06	3.963e-07	4.984	8.09e-07 ***
Electorate	-1.453e-06	2.722e-07	-5.339	1.31e-07 ***
econInactiveRetiredPcnt	1.318e-02	1.139e-03	11.570	< 2e-16 ***

Signif. codes:
0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
Residual standard error: 0.08802on 623 degrees of freedom
Multiple R-squared: 0.4112, Adjusted R-squared: 0.4055
F-statistic: 72.52 on 6 and 623 DF, *p*-value: < 2.2e-16

When combined with the political value space mapped with Twitter messages, the model that incorporates demographic data accounts for nearly half of the variance found in the referendum results (R^2_{adj}=.41, p<2.2e-16). Table 1 shows that this model improves the previous one, which lacked demographic data, thus foregrounding the possibilities of complex social data modeling by mixing social media signal with demographic data that can be aggregated at user, group, or community levels. Although social media data remain a non-representative sample of the larger population, they can provide important markers for understanding the evolution of public debates and the geographic coverage of the discussion (Bastos, Recuero, & Zago, 2014). The results of the classifier also shed light on the importance of economic issues that might have been of vital importance to the user base tweeting the referendum, a component of the Brexit debate overshadowed by the much-discussed cleavage between the metropolitan elite in London and parts of England and Wales that were economically worse off.

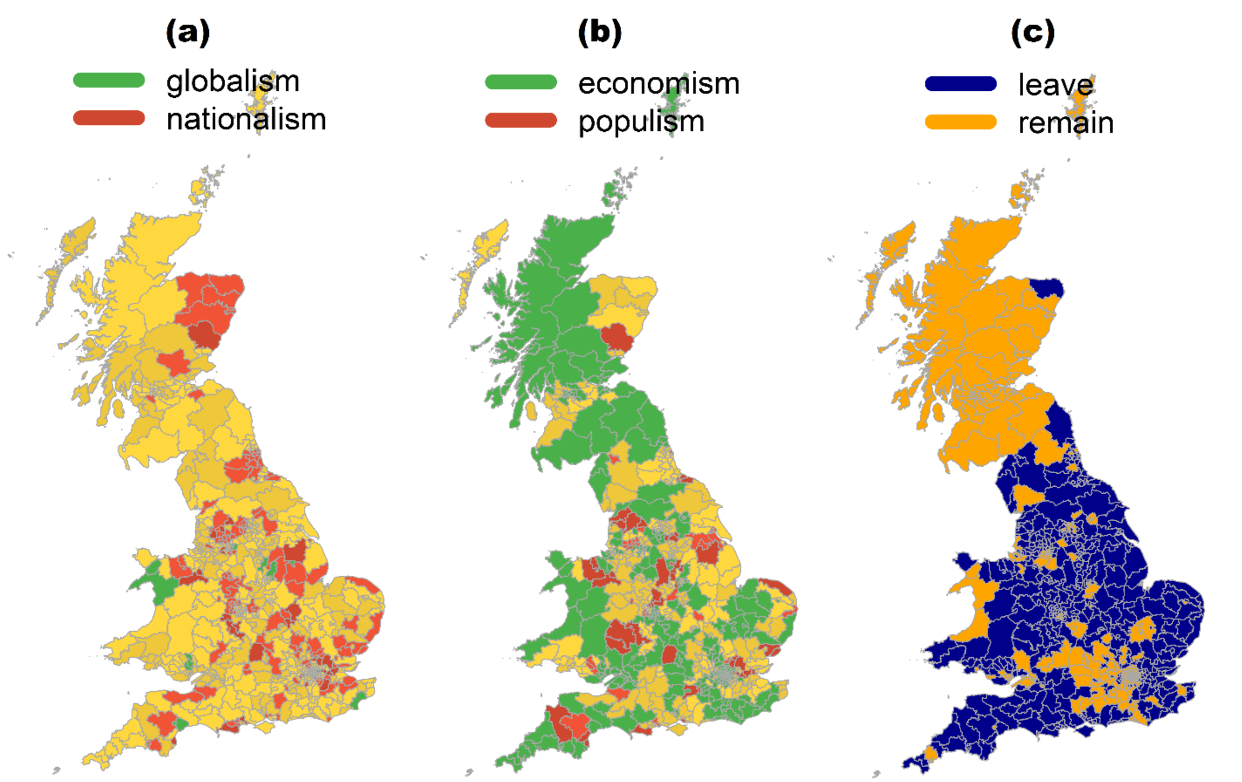

Figure 4: Mean score of Globalism-Nationalism (a) and Economism-Populism (b) for each Parliamentary Constituency and the results of the referendum (c)

However, these results only partially support hypothesis H4—that the ideological value space can be used to model public opinion formation relative to the referendum results. More than half of the variance found in the referendum results remain unaccounted by the model and a closer inspection of the aggregate scores for Globalism, Nationalism, Economism, and Populism show that the map only partially matches the results of the referendum (Figure 4c). Apart from London and north-west Wales (Gwynedd), globalist messages are absent in Figure 4a, with nationalist content appearing in Scotland (which voted Remain, but has long contended with a nationalistic agenda pressing for an independent Scotland), the English Midlands, and the north of England. Populist messages are also relatively underwhelming covering only portions of the Midlands and North (Figure 4b). It is the economic discourse that is prevalent in the debate registered in the Twittersphere, being particularly prominent in Scotland, north-west Wales, and Greater London (Figure 4b). As an expression of the public opinion, Twitter debate appears invariably focused on economic and nationalistic issues as opposed to the populist and globalist sentiments thought to have shaped much of the referendum campaign (Inglehart & Norris, 2016).

Discussion and Conclusion

The results reported in this study have partially upset our expectations. We did not find, for one, that economically fragile northern Britain was any more likely to embrace nationalist content. In fact, it was Scotland that appeared as a relatively fertile ground for nationalist messages. Furthermore, we rejected hypothesis H1 which posited that the distribution of users using hashtags to advocate either side of the campaign would mirror the results of the referendum.

These results contradict previous research that found the expression of agreement with a topic on Twitter to predict the results of the referendum (Celli, et al., 2016)—at least at more granular levels such as Parliamentary Constituencies. Although hypotheses H2 and H3 are partially supported, the distribution of globalist, nationalist, populist, and economist content is somewhat at odds with the geographic distribution of the Leave-Remain vote. More importantly, the significant results reported with the model tested for hypothesis H4 relies heavily on the AIC stepwise model selection that incorporated variables exogenous to the Twitter network such as unemployment, valid votes, electorate, and ratio of retired population living in the constituency.

Twitter conversation in the weeks leading up to the referendum vote was largely centred on nationalistic and economic sentiments, a result that sheds light on the central research question investigated in this study—i.e., that not only material inequality, but also ideological realignments have contributed to the outcome of the referendum. On the one hand, the variables that have improved the model are associated with issues surrounding material inequality, chief of which are the percentage of economically active residents and the size of the parliamentary constituency. On the other hand, ideological orientation has also proven capable of explaining the unexpected outcome of the U.K. public to leave the E.U. As such, our results suggest that it was primarily outrage at material inequality along with a nationalistic upsurge that can help explain this epochal change in British politics (Inglehart & Norris, 2016), a result somewhat at odds with literature foregrounding the resurgence of populism in Western industrialized countries (Tollefson, 2016).

In summary, we found evidence that nationalism was a quintessential component of the referendum debate during most of the campaign, with three quarters of messages having some degree of nationalistic sentiment embedded in them. These results however need to be considered in the context of the limited heuristic value of this ideological coordinate. While the classifier successfully identifies nationalistic sentiments in Scotland and the Midlands, these areas have voiced fundamentally different versions of nationalism. In other words, although nationalism appears to have been a critical marker of the Brexit value space, there are important differences in Scottish and English nationalism that extrapolate the heuristic confines of the classifier. Populist messages, however, were decidedly of lesser importance compared with the sheer volume of tweets discussing the economic consequences of Britain leaving the E.U., a trend that is however inverted as we approach the date of the vote.

With this paper, we aimed to advance a proof of concept research design employing social media signal to model the ideological value underpinning the Brexit referendum outcome. The continuous streaming of Twitter messages can be leveraged to identify short and long-term shifts that are difficult to detect with surveys and interview instruments. As such, the rationale of the methodology advanced in this study is to employ social media data as geographically-rich intelligence that can be explored and combined with established social science research methods. The strength of the approach explored herewith lies less in its predictive or forecasting power (Jungherr & Jürgens, 2013; Jungherr, et al., 2012; Lazer, Pentland, Adamic, Aral, Barabási, Brewer, Christakis, Contractor, Fowler, Gutmann, Jebara, King, Macy, Roy, & Van Alstyne, 2009) and more on the range of possibilities for exploring ongoing developments that would otherwise require the extensive, continuous, and expensive use of survey methodologies.

Our analysis is notably restricted to the period of the campaign, with no insights as to how the debate evolved following the referendum. The data were also aggregated both at the geographic level (Parliamentary Constituencies) and on the temporal scale (no longitudinal variance was measured or incorporated into the model). Future research should explore the temporal variation of ideological coordinates that can be continuously mined and detected with a machine learning algorithm such as the one trained for this study. While the fitted model presented herewith has limited explanatory power relative to the outcome of the referendum, greater variability in the opinions of users tweeting campaigns in the run-up to the vote could potentially be detected with such an approach. The temporal component of social media data could also be integrated into the model to explore relationships that might not have been possible to measure otherwise. In short, the results of this study present important insights into the reasons why the British public decided to end the country's membership of the E.U., but future research should seek to further investigate the sentiments and ideological positioning that might have crystallized through the referendum debate.

References

Anstead, N., & O'Loughlin, B. (2015). Social Media Analysis and Public Opinion: The 2010 UK General Election. *Journal of Computer-Mediated Communication, 20*(2), 204-220. doi: 10.1111/jcc4.12102

Asthana, A., Quinn, B., & Mason, R. (2016). UK votes to leave EU after dramatic night divides nation, *The Guardian*. Retrieved from https://www.theguardian.com/politics/2016/jun/24/britain-votes-for-brexit-eu-referendum-david-cameron

Bakshy, E., Messing, S., & Adamic, L. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science, 348*(6239), 1130-1132. doi: 10.1126/science.aaa1160

Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting From Left to Right: Is online political communication more than an echo chamber? *Psychological Science, 26*(10), 1531-1542. doi: 10.1177/0956797615594620

Bastos, M. T., & Mercea, D. (2017). The Brexit Botnet and User-Generated Hyperpartisan News. *Social Science Computer Review*. doi: 10.1177/0894439317734157

Bastos, M. T., Recuero, R. C., & Zago, G. S. (2014). Taking tweets to the streets: A spatial analysis of the Vinegar Protests in Brazil. *First Monday, 19*(3). doi: 10.5210/fm.v19i3.5227

Bauman, Z. (2012). *Liquid modernity*. Cambridge: Polity.

BBC. (2016). EU Referendum: The result in maps and charts, *BBC News*. Retrieved from http://www.bbc.co.uk/news/uk-politics-36616028

Becker, S., O., Fetzer, T., & Novy, D. (2016). Who voted for Brexit? A Comprehensive District-Level Analysis (pp. 1-69). Coventry: Centre for Competitive Advantage in the Global Economy, University of Warwick.

Celli, F., Stepanov, E. A., Poesio, M., & Riccardi, G. (2016). *Predicting Brexit: Classifying Agreement is Better than Sentiment and Pollsters.* Paper presented at the Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media (PEOPLES), Osaka, Japan.

Crouch, C. (1997). The Terms of the Neo-Liberal Consensus. *The Political Quarterly, 68*(4), 352-360.

Cummings, D. (2016). On the Referendum #20: The Campaign, Physics, and Data Science - Vote Leave's 'Voter Intention Collection System' (VICS) now available for all. Retrieved from https://dominiccummings.wordpress.com/2016/10/29/on-the-referendum-

20-the-campaign-physics-and-data-science-vote-leaves-voter-intention-collection-system-vics-now-available-for-all/

Dahlgren, P. (2009). *Media and Political Engagement: Citizen, Communication and Press*. Cambridge: Cambridge University Press.

Dunleavy, P., & Margetts, H. (2001). From Majoritarian to Pluralist Democracy? *Journal of Theoretical Politics, 13*(3), 295-319. doi: doi:10.1177/095169280101300304

Dutton, W. H., Reisdorf, B. C., Dubois, E., & Blank, G. (2017). A Cross-National Survey of Search and Politics (March 31, 2017). Quello Center Working Paper No. 2944191 Retrieved 4 April 2017, from https://ssrn.com/abstract=2944191

Fletcher, R., & Nielsen, R. K. (2017, June 21, 2017). Using social media appears to diversify your news diet, not narrow it. *NiemanLab*.

Habermas, J. (1991). *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Cambridge: MIT Press.

Hanretty, C. (2017). Areal interpolation and the UK's referendum on EU membership. *Journal of Elections, Public Opinion and Parties*, 1-18.

Huyen Do, V., Thomas-Agnan, C., & Vanhems, A. (2015). Spatial reallocation of areal data – another look at basic methods. *Revue d'Économie Régionale & Urbaine, mai*(1), 58. doi: 10.3917/reru.151.0027

Inglehart, R. F., & Norris, P. (2016). *Trump, Brexit, and the Rise of Populism: Economic Have-nots and Cultural Backlash*. Paper presented at the American Political Science Association Annual Meeting, Philadelphia, USA. https://research.hks.harvard.edu/publications/getFile.aspx?Id=1401

Jungherr, A., & Jürgens, P. (2013, October 23-26). *Forecasting the pulse: how deviations from regular patterns in online data can identify offline phenomena.* Paper presented at the Internet Research 14.0, Denver, USA.

Jungherr, A., Jürgens, P., & Schoen, H. (2012). Why the Pirate Party Won the German Election of 2009 or The Trouble With Predictions: A Response to Tumasjan, A., Sprenger, T. O., Sander, P. G., & Welpe, I. M. “Predicting Elections With Twitter: What 140 Characters Reveal About Political Sentiment”. *Social Science Computer Review, 30*(2), 229-234. doi: 10.1177/0894439311404119

Kriesi, H., & Frey, T. (2008). The United Kingdom: moving parties in a stable configuration. In E. Grande, H. Kriesi, M. Dolezal, R. Lachat, S. Bornschier & T. Frey (Eds.), *West European Politics in the Age of Globalization* (pp. 183-207). Cambridge: Cambridge University Press.

Kriesi, H., Grande, E., Lachat, R., Dolezal, M., Bornschier, S., & Frey, T. (2008). Globalization and its impact on national spaces of competition. In E. Grande, H. Kriesi, M. Dolezal, R. Lachat, S. Bornschier & T. Frey (Eds.), *West European Politics in the Age of Globalization* (pp. 3-22). Cambridge: Cambridge University Press.

Laclau, E. (2005). *On populist reason*. London: Verso.

Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., . . . Van Alstyne, M. (2009). Computational Social Science. *Science, 323*(5915), 721-723. doi: 10.1126/science.1167742

McLaren, L. M. (2002). Public Support for the European Union: Cost/Benefit Analysis or Perceived Cultural Threat? *The Journal of Politics, 64*(2), 551-566.

Mudde, C. (2000). *The ideology of the extreme right*. Manchester: Manchester University Press.

Mudde, C. (2004). The Populist Zeitgeist. *Government and Opposition, 39*(4), 541-563. doi: 10.1111/j.1477-7053.2004.00135.x

Mudde, C., & Rovira Kaltwasser, C. (2012). *Populism in Europe and the Americas* (Vol. Cambridge): Cambridge University Press.

Nilsson, B., & Carlsson, E. (2014). Swedish politicians and new media: Democracy, identity and populism in a digital discourse. *New Media & Society, 16*(4), 655-671. doi: 10.1177/1461444813487964

ONS. (2011). *2011 Census*. Retrieved from https://www.ons.gov.uk/census/2011census.

ONS Geography. (2011). National Statistics Postcode Lookup UK. In O. f. N. Statistics (Ed.). London: ONS.

ONS Geography. (2017). *NHS Postcode Directory User Guide*. Retrieved from https://digital.nhs.uk/media/553/User-guide/pdf/user_guide.

Parker, G. (2016, 12/12/2016). Nigel Farage: Eurosceptic scourge of the ‘political elite’, *Financial Times*. Retrieved from https://www.ft.com/content/ab8b3b98-be3a-11e6-8b45-b8b81dd5d080

Pew Research Center. (2013). The Demographics of Social Media Users, 2012. In M. Duggan & J. Brenner (Eds.). Washington, D.C.: Pew Research Center's Internet & American Life Project.

Piketty, T. (2014). *Capital in the twenty-first century*. Cambridge Massachusetts: The Belknap Press of Harvard University Press.

Rennie Short, J. (2016). The Geography of Brexit: What the Vote Reveals about Disunited Kingdom, *The Conversation*. Retrieved from http://theconversation.com/the-geography-of-brexit-what-the-vote-reveals-about-the-disunited-kingdom-61633

Selivanov, D. (2016). text2vec: Modern Text Mining Framework for R (Version 0.4.0): CRAN. Retrieved from https://CRAN.R-project.org/package=text2vec

Sunstein, C. R. (2007). *Republic.com 2.0*. Princeton, N.J: Princeton University Press.

Tollefson, J. (2016). Researchers baffled by nationalist surge. *Nature, 540*(7632), 182–183. doi: 10.1038/540182a

Tumasjan, A., Sprenger, T., Sandner, P., & Welpe, I. (2011). Election Forecasts With Twitter: How 140 Characters Reflect the Political Landscape. *Social Science Computer Review*. doi: citeulike-article-id:8506921

Turner, B. (2002). *Orientalism, postmodernism and globalism*. Singapore: Routledge.

Wills, J. (2015). Populism, localism and the geography of democracy. *Geoforum, 62*, 188-189. doi: http://dx.doi.org/10.1016/j.geoforum.2015.04.006

Woods, D. (2009). Pockets of resistance to globalization: the case of the Lega Nord. *Patterns of Prejudice, 43*(2), 161-177.