

What Drives Anti-Immigrant Sentiments Online? A Novel Approach Using Twitter

Anastasia Menshikova^{1,*} and Frank van Tubergen^{2,3}

¹Institute for Analytical Sociology, Linköping University, Norrköping, 60174, Sweden, ²Department of Sociology, Utrecht University, 3584 CH Utrecht, Netherlands and ³Netherlands Interdisciplinary Demographic Institute (NIDI), KNAW/University of Groningen, 2511 CV The Hague, Netherlands

*Corresponding author. Email: anastasia.menshikova@liu.se

Submitted July 2020; revised December 2021; accepted January 2022

Abstract

Most studies use survey data to study people's prejudiced views. In a digitally connected world, research is needed on out-group sentiments expressed online. In this study, we show how one can elaborate on existing sociological theories (i.e. group threat theory, contact theory) to test whether anti-immigrant sentiments expressed on Twitter are related to sociological conditions. We introduce and illustrate a new method of collecting data on online sentiments, creating a panel of 28,000 Twitter users in 39 regions in the United Kingdom. We apply automated text analysis to quantify anti-immigrant sentiments of 500,000 tweets over a 1-year period. In line with group threat theory, we find that people tweet more negatively about immigrants in periods following more salient coverage of immigration in the news. We find this association both for national news coverage, and for the salience of immigration in the personalized set of outlets people follow on Twitter. In support of contact theory, we find evidence to suggest that Twitter users living in areas with more non-western immigrants, and those who follow a more ethnically diverse group of people, tweet less negatively about immigrants.

Introduction

Over the last decades, Europe has experienced multiple immigration waves, which have strongly increased the population with an immigrant background. A large share of native Europeans express opposition towards immigration and show negative attitudes towards ethnic minority members (Meuleman, Davidov and Billiet, 2009). As a response to the widespread concern about polarization, discrimination, and ethnic exclusionary tendencies in society, a substantial body of research has emerged on anti-immigrant attitudes in European societies (Ceobanu and Escandell, 2010).

The common approach to study anti-immigrant attitudes is to use a survey and ask respondents to anonymously indicate their opinion about foreigners on a standardized scale. This line of research has documented important trends (Semyonov, Raijman and Gorodzeisky, 2006; Bohman and Hjerem, 2016) and cross-country differences (Scheepers, Gijsberts and Coenders, 2002; Savelkoul et al., 2012) in anti-immigrant attitudes. It has also provided a wealth of empirical findings about the role of intergroup contact (Allport, 1954; Dovidio, Gaertner and Kawakami, 2003) and group threat (Blumer, 1958; Blalock, 1967) in shaping intergroup

attitudes. For example, studies have examined the impact of immigration flows and the size of ethnic minority groups (Pottie-Sherman and Wilkes, 2017), the influence of salience and framing of immigration in the news (Boomgaarden and Vliementhart, 2009; Schlueter and Davidov, 2013; van Klingeren et al., 2015; Czymara and Dochow, 2018; Eberl et al., 2018), and the effect of economic conditions (Weber, 2015; Gorodzeisky and Semyonov, 2016; Kuntz, Davidov and Semyonov, 2017; Meuleman et al., 2020).

The world has changed, however, and with the rise of social media and digital communication, people can also express their anti-immigrant opinions publicly, such as on Twitter and Facebook. The few studies that have been done in this area so far show that negative out-group sentiments and hate speech on social media platforms are pervasive and on the rise, but also strongly fluctuate over time (Williams, 2019). One study observed that negative out-group sentiments on social media peaked in the aftermath of the Woolwich, London terrorist attack in 2013 (Williams and Burnap, 2016). Another study found evidence for a ‘celebrity exposure effect’: Liverpool F.C. fans halved their rates of posting anti-Muslim tweets on Twitter relative to fans of other top-flight English soccer clubs, after the rise to fame of Liverpool F.C. soccer star Mohamed Salah—a player whose Muslim identity is well-known among the supporters (Arababa’h et al., 2021). However, most studies in this field use a qualitative case-study approach (Bliuc et al., 2018) and little is known about the sociological drivers of online negative out-group sentiments and hate speech (Flores, 2017).

A key question for sociological research on interethnic attitudes is whether existing theories and stylized empirical patterns can be used for understanding these online expressions of intergroup sentiments or that one should develop an entirely new theoretical framework. In defence of the latter position, one may argue that those who express their views on online platforms, such as on Twitter, are a selective group (Mellon and Prosser, 2017). In the United Kingdom, for example, data for 2013 show that only 30 per cent of the population used Twitter; younger people and those with higher income being overrepresented (Blank and Lutz, 2017). In addition, one could argue that, in line with the ‘preference falsification’ mechanism (Kuran, 1997), people’s publicly revealed online sentiments may not reflect their private opinions (Flores, 2017). Furthermore, on social media platforms, users are often exposed to a selective group of like-minded people (Bakshy, Messing and Adamic 2015), which may affect people’s publicly expressed sentiments as well.

To date, however, little is known whether, despite these particular conditions, hypotheses derived from existing sociological theories work well for predicting online intergroup sentiments. If they do, it would mean that patterns are robust for selective samples and preference falsification, and that public opinions expressed on social media, such as on Twitter, are not just noise or subject to entirely different processes for which one needs to develop a new theoretical framework.

In this study, we argue that online data sources provide a rich and free tool for sociologists to monitor and study intergroup sentiments. We introduce a new method of data collection, in which one observes the same Twitter users over time. We created a panel of more than 28,000 Twitter users across 39 regions in the UK, which we followed for a 1-year period (1 November 2018 to 31 October 2019). We studied their daily Twitter posts and, using automated text analysis, quantified their sentiments towards immigrant out-groups. We show how one can elaborate on existing sociological theories (i.e. group threat theory, contact theory) to test whether online anti-immigrant sentiments expressed on Twitter are related to sociological conditions.

Theory and Hypotheses

Media Salience

We examine whether fluctuations in online negative sentiments towards immigrants are related to changes in the salience of immigration in the news. Media and news messages can set the agenda for topics which people discuss (Scheufele and Tewksbury, 2007), and also shape people’s (mis)perceptions. Based on group threat theory, a key hypothesis is that the salience of the topic of immigration in national news outlets increases anti-immigrant sentiments (Schemer, 2012; Van Klingeren et al., 2015). Frequent mentioning of topics like ‘immigrants’, ‘immigration’, and ‘refugees’ in the news may create the perception of a large, or increasingly large, group of immigrants in society. Such perceptions may then trigger feelings of threat, as ethnic majority members consider immigrants as competitors for important resources (Blumer, 1958; Blalock, 1967). They may fear that immigrants take away jobs or undermine mainstream cultural norms and values. Hence, it is argued that frequent exposure to media messages on immigration issues amplifies anti-immigrant sentiments. Previous studies have found ample support for this hypothesis (Czymara and Dochow, 2018; Eberl et al.,

2018; Berning, Lubbers and Schlueter, 2019; Boer and van Tubergen, 2019).

We extend this literature by examining whether online negative sentiments, as expressed on Twitter, are also fluctuating with the amount of news on immigration. Although public expressions on Twitter may be different from people's private views, and those who tweet are not representative of the larger population, people's sentiments expressed in their tweets may nevertheless be associated with *changes* in the visibility of immigration in the news. We therefore hypothesize:

H1a: The more strongly Twitter users are exposed to national news on immigration topics, the more negative their sentiments towards immigrants expressed on Twitter become.

Nowadays, people use multiple online sources to read the news. The news landscape has become increasingly personalized and fragmented, which means that on the same day, two persons may be differently exposed to news about immigration. It may be highly salient in the news consumed by one, while the other did not read anything about immigration. We examine which news outlets people follow on Twitter, and whether the salience of immigration in their personalized news impacts their sentiments expressed online:

H1b: The more strongly Twitter users are exposed to news on immigration topics via the news outlets they follow on Twitter, the more negative their sentiments towards immigrants expressed on Twitter become.

Non-Western Group Size in a Region

The actual size of the immigrant group in people's region of living may also be a driver of anti-immigrant sentiments. Immigrant group size shapes people's personal experiences, and that of the people they know personally, which is another source of information about outgroups. The way group size affects people's experiences is debated, however, and two (opposing) forces have been proposed (Schlueter and Scheepers, 2010).

Group threat theory, on the one hand, emphasizes that people feel more threatened by immigrants when they live in a region with many immigrants. Competition for jobs is then more intense, and people may fear for their safety as well, as they may believe that immigrants are more criminal and aggressive than ethnic majority members. In particular, non-western immigrants are seen as threatening, because of differences in religion, norms, and values. Group threat theory therefore expects that non-western immigrant group size in

the region of living is associated with more-negative sentiments towards immigrants.

Contact theory, on the other hand, argues that personal experiences with immigrants lead to more positive feelings about immigrants (Allport, 1954; Dovidio, Gaertner and Kawakami, 2003; Pettigrew and Tropp, 2008). Because immigrant group size is positively correlated with the opportunities for intergroup contacts, it follows that immigrant group size in the region of living may lead to less-negative sentiments towards immigrants. The empirical evidence for the impact of immigrant group size is still a matter of discussion (Pottie-Sherman and Wilkes, 2017). We take this discussion in a new direction by looking at sentiments expressed on Twitter:

H2a: The higher the proportion of non-western immigrants within the region of residence of Twitter users, the more negative are their sentiments towards immigrants expressed on Twitter (group threat theory).

H2b: The higher the proportion of non-western immigrants within the region of residence of Twitter users, the less negative are their sentiments towards immigrants expressed on Twitter (contact theory).

The presence of a larger share of (non-western) immigrants in the region of living may not necessarily imply that ethnic majority members have actual intergroup contacts. Foci, such as schools and workplaces, can be highly segregated. It is therefore important to examine actual intergroup contacts. Such cross-group contacts may be the result of selection processes, but research suggests that they are also, and even more so affecting people's attitudes (Pettigrew and Tropp, 2008). We take this literature in a new direction by studying people's Twitter connections to members of non-western immigrant groups. Being exposed to and interacting with a more-diverse group of people online can lead to a reduction of prejudice and out-group negativity. We therefore expect:

H2c: The higher the proportion of non-western immigrants among the connections of Twitter users, the less negative are their sentiments towards immigrants expressed on Twitter (contact theory).

Media Salience and Group Size in a Region

Most studies have focused on how media messages impact negative attitudes towards immigrants, or how group size affects attitudes, without reflecting on the possible interactions between these two social forces. Schlueter and Davidov (2013) were among the first to

study the interplay between media exposure and group size, finding that the salience of immigration in the news has a particularly negative effect in areas with fewer immigrants. Subsequent studies have equally found evidence for this relationship (Czymara and Dochow, 2018). One possible mechanism playing a role is that people rely more strongly on the news to gather information about outgroups when they have fewer personal experiences with those outgroups (Zucker, 1978). Following contact theory, intergroup encounters—which are more common in areas with more immigrants—thus provide a strong correction to media messages and reduce the need to rely on the news to develop outgroup. Therefore, we test the following hypothesis:

H3: The higher the proportion of non-western immigrants within the region of Twitter users, the weaker the effect of immigration-related news on their anti-immigrant sentiment expressed on Twitter.

Data and Methods

Twitter

Twitter is a widely used microblogging platform. On Twitter, users can publish short posts (tweets) up to 280 characters, reply to tweets posted by other users, or re-post content from other Twitter pages. Users can provide brief information about themselves, including first and last name, date of birth, geolocation, and a description in a free form. Users can follow each other and read the content posted by followed pages in their Twitter feed. Twitter allows collecting publicly available data via Application Programming Interface (hereafter as Twitter API) and Twitter Advanced Search engine, which we use in this study employing R and Python packages.

Data Collection

Our data collection from Twitter resembles the design of a panel survey. Over a year, we followed the same panel of Twitter users in the United Kingdom and collected all their Twitter posts related to immigration and foreign outgroups. To clarify, we define foreign outgroups as groups that have a migration history in the UK context. It thus means that we study sentiments towards a broad category of foreign outgroups that includes immigrants, refugees, asylum seekers, ethnic minorities, Muslims, and also specific ethnic groups, such as Poles, Chinese, and Syrians.

To identify tweets related to immigration and foreign outgroups, we developed a search string that includes

keywords capturing these categories (see [Supplement for details](#)). First, we included general terms frequently used in discourse with regard to immigration in the UK context, such as ‘immigration’, ‘refugee’, ‘foreigner’, and ‘Muslim’ (Schlueter and Davidov, 2013; Bleich et al., 2015; McLaren, Boomgaarden and Vliegthart, 2018). Second, we added the 15 largest immigrant groups living in the United Kingdom (Office for National Statistics, 2018) (e.g. ‘Africans’, ‘Poles’, ‘Pakistani’), as well as the main refugee groups of the 2015 refugee crisis (e.g. ‘Syrian’, ‘Afghani’, ‘Iraqi’). We did not include ‘Italian’, ‘German’, ‘French’, and ‘Irish’ as keywords, even though the size of corresponding national groups living on the territory of United Kingdom is prominent. We did so because (i) they are highly skilled and seen as less-threatening (Hainmueller and Hiscox, 2010; Schlueter and Davidov, 2013) and (ii) these terms return many irrelevant tweets (for instance, about Italian cuisine, etc.). For each of the terms from the list, the automated calls to Twitter Advanced Search engine were composed in a way that it returned tweets with all the word forms of a seed word, i.e. both singular and plural. In order to avoid language misspecifications, two British-born native speakers independently reviewed the final set of search terms.

Next, we developed a list of locations in the United Kingdom from which we sampled the Twitter users. Sampling by locations was done to test the hypotheses about the regional share of the non-western population. We used NUTS classification (Nomenclature of Territorial Units for Statistics established by Eurostat) to define geographical units of the United Kingdom. NUTS is a hierarchical geocode system, and we used level-2 for the sampling and main analysis of the study. Within each of the 40 level-2 regions, we selected the most populous cities or towns (population greater than 50,000) and used the names of those 183 places to sample Twitter users across the United Kingdom (see [Supplement](#)).

We then created combinations of every search term (related to the topic of immigration) with every location (i.e. 37 search terms \times 183 places = 6,771). Subsequently, we sent calls to the Twitter Advanced Search engine using the `GetOldTweets3` package in Python (Mortl, 2018) to gather all immigration-related tweets from the list of locations predefined, for 1 year. The period for which we collected data was 1 November 2018 to 31 October 2019. We obtained a corpus of 1,143,818 tweets posted by 44,594 users.

Not all users and tweets were useful to keep in our panel, however. We documented this filtering process for reasons of transparency and reproducibility (see

Supplement). We automatically filtered out users based on the following criteria: (i) the user posted less than two tweets during the year; (ii) all the tweets of the user were posted within 1 day (iii); and the location of the user returned by the Twitter API was non-identifiable or irrelevant. In order to automatically filter out organizations, celebrities, or other influential people, (iv) we excluded users who follow more than 5,000 Twitter pages, and users with more than 9,361 followers (i.e. highest 10 per cent). Next, we excluded irrelevant posts, i.e. tweets that are not related to immigrant groups in the United Kingdom, but to other topics, such as about issues abroad (e.g. the war in Syria, Muslims in Indonesia, Mexican immigrants in the United States). In the end, we arrived at a panel of 28,161 Twitter users in the United Kingdom, who together tweeted 534,955 posts about immigration in 1-year time.

Measurements

Anti-immigrant sentiments

To measure anti-immigrant sentiments expressed on Twitter among the Twitter users from our panel, we used automated sentiment analysis. Sentiment techniques have been widely used for extracting sentiment polarity of user-generated content on Twitter, such as to study attitudes towards climate change (Cody et al., 2015). Employing sentiment analysis for studying attitudes towards immigrants has been uncommon so far (an exception is Flores, 2017). We used a lexicon-based approach in our study. Lexicon-based sentiment analysis assigns a sentiment score to a piece of text based on the polarity scores of words that are contained in the text. These polarity scores are developed by experts, who have made extensive lists of words and corresponding sentiments. For instance, the word ‘wonderful’ might have a score +3 and the word ‘disgusting’ might have a score -2.

Sentiment detection was performed with SentiStrength (Thelwall, 2017). In a recent benchmark comparison of state-of-the-practice sentiment analysis, SentiStrength came out as the ‘winner’ in almost all contexts (Ribeiro et al., 2016). It particularly performs well for Twitter, with an accuracy of 95 per cent or higher, when it was tested on datasets that were human coded (Ribeiro et al., 2016). A strong feature of SentiStrength is that the positive (or negative sentiment) of a lexicon of 2,310 sentiment words and word stems were human assigned (Thelwall, 2017). That is also the case for emoticons, exclamation marks, capital letters, etc.—this is based on human coding and extensive validation checks.

SentiStrength assigns each tweet with negative (ranging from -5 to -1) and positive (from 1 to 5) scores, based on the words in a tweet recognized by the algorithm as positive or negative. We created a general sentiment score for every tweet by adding positive and negative scores together. We then reversed the sentiment scale, as our dependent variable is anti-immigrant attitudes. In this way, we have a variable scaled from -4 to 4 where higher values denote stronger anti-immigrant sentiments. We used the Python 3 wrapper for SentiStrength (Hung, 2020).

An alternative to the (unsupervised) lexicon approach is supervised machine learning. With this method, one codes a set of tweets (by at least two independent human coders) in terms of their sentiments towards immigrants. This dataset is then used to train an algorithm, which can then be applied to the entire corpus of tweets. We refrained from supervised machine learning, for three reasons: (i) the sample of tweets that needs to be coded is very large to construct a valid test data set, (ii) transparency and replication become less-straightforward when refining algorithms with additional supervised learning, and (iii) studies show that the lexicon-based unsupervised mode of SentiStrength has similar overall accuracy to that of supervised modes (Thelwall, 2017).

Salience of immigration in the national news

To measure the salience of immigration in the news, scholars have typically looked at news in prominent newspapers as a proxy for the national news climate (Dunaway et al., 2011; Schlueter and Davidov, 2013; Eberl et al., 2018; McLaren, Boomgaarden and Vliegenthart, 2018). We used *The Guardian* and *The Sun* as two major news outlets in the United Kingdom, and collected the tweets from their accounts for the period of study (1 November 2018 to 31 October 2019). To detect immigration-related news tweets, we used the same list of words which we used for collecting users’ tweets on immigration (see above). We then created the variable immigration-related news as the average percentage of news mentioning immigrant-related topics.

It should be emphasized that current theory does not specify which time lag should be used. Previous empirical work has used different specifications. Some studies aggregated the media salience of immigration for a period of half a year (van Klinger et al., 2015), others used monthly figures (Boomgaarden and Vliegenthart, 2009; Schemer, 2012), 3 weeks (Czymara and Dochow, 2018), weekly data (Dunaway et al., 2011), 1-day lags (Boer and van Tubergen, 2019), or immediate effects in

experimental designs (Jacobs and van der Linden, 2018). In the main analyses, we use the average percentage of news mentioning immigrant-related topics in the week before, and present results from additional analyses using different time lags.

Salience of immigration in personalized news

We used the news sources to which each user in our panel subscribed on Twitter to develop a personalized measure of users' weekly exposure to immigration-related news. We gathered the news subscriptions of users from our sample via accessing Twitter API with R package 'rtweet' (Kearney, 2019). Across all collected subscriptions of individuals on Twitter, we detected the category of 'news pages' in the following way. First, we compiled a list of the 50 most popular media sources in the United Kingdom and collected usernames of their official pages on Twitter.¹ We then detected these news pages among Twitter subscriptions of the individuals from our sample. However, it appeared that users follow other news sources on Twitter, such as international media outlets or media from other countries (e.g. CNN, Fox News), and also regional newspapers. Therefore, we expanded our list of news sources. First, across the list of Twitter subscriptions, we selected sources followed by more than 500 users from our sample. Among those popular sources, we checked every account for whether it was indeed a news outlet, based on (i) the verified status on Twitter and (ii) the corresponding website. In this way, we added 90 news outlets to our list, which makes up 160 Twitter news pages in total (see Supplement). We collected all the tweets posted by the 160 news outlets for the same period as we collected user tweets for (1 November 2018 to 31 October 2019).

Then, we calculated a percentage of immigration-related tweets posted by each news source on a weekly basis. Because many users follow more than one news page on Twitter, we calculated the average salience of immigration across the news pages they are following on Twitter to capture their salience of immigration-related news. Therefore, for every user, we obtained an average weekly media exposure to immigration-related news. Additionally, we also explore different time-lags to examine the effects of shorter or longer periods of media exposure (e.g. 1 day, 1 week, 1 month).²

Non-western immigrants in region

We measured the percentage of the population that is non-western immigrant within the region in which the Twitter users from our panel reside. We used figures for the NUTS-2 regions in the United Kingdom for 2018 (Office for National Statistics, 2018). We then identified

geo-locations of the Twitter users in our panel, matching their place of living indicated on Twitter with the corresponding NUTS-2 region-level data.

Non-western connections on Twitter

We included a proxy of the share of non-western immigrant ties on Twitter, as captured by outgoing links (i.e. whom a user follows). For this, we used account names and apply the ethnicity-recognition algorithm 'ethnicolr' in Python (Sood and Laohaprapanon, 2018). Ethnicolr detects ethnic origin based on the sequence of letters in the first and last name; it returns the list of probabilities of belonging to 1 of 13 ethnic groups (see Supplement). We employed the version of the algorithm trained on Wikipedia, which results in predictions with an average precision score of 0.73.

We counted the following regions as non-western: Asia, Africa, Eastern Europe. We assigned ethnicity based on the highest probability predicted, with the probability of belonging to a certain ethnic group no less than 0.5, otherwise we removed the cases because of high uncertainty. For every user in our sample, we calculated the percentage of non-western immigrant connections as a ratio of connections recognized by the algorithm as non-western immigrants to all the connections with successfully predicted ethnicity.

Control variables

We controlled for *population density* in the region of living (NUTS-2), i.e. the number of people residing per square kilometre (Eurostat, 2018). We also controlled for the (log) *number of followers* and the (log) *number of connections* on Twitter.

Table 1 shows descriptive statistics for the main variables we use in the analysis.

Analytical strategy

We ran multilevel linear regressions to test the hypotheses. As we conducted the analysis on a daily level, we aggregated the more than 500,000 tweets posted to the average sentiments for each user per day. We then have a three-level multilevel structure, in which time-varying daily data on the 314,589 observations of sentiments for each user (level 1) are nested within 28,161 Twitter users (level 2), and these Twitter users are nested within the 39 NUTS-2 regions (level 3). All models presented were estimated with maximum likelihood estimation. We conducted the modelling using R package 'lmerTest' for mixed models (Kuznetsova, Brockhoff and Christensen, 2017).

Table 1. Descriptive statistics

	Min	Max	Mean	SD
Anti-immigrant sentiments on Twitter (SentiStrength)	-4.00	4.00	-0.19	1.31
Saliency immigration national news, 1 week	0.22	6.45	2.11	0.86
Saliency immigration personalized news, 1 week	0.00	52.97	3.61	3.00
Non-western immigrants in region (per cent)	2.69	34.76	16.34	10.63
Non-western ties on Twitter (per cent)	0.00	100.00	42.82	18.57

Results

Findings from the null-model (not presented here) reveal that 81.82 per cent of the variance in anti-immigrant attitudes expressed on Twitter is at the user-daily level, 17.97 per cent at the user level and 0.21 per cent at the region level. This suggests that anti-immigrant sentiments expressed on Twitter strongly fluctuate over time, significantly differ across Twitter users, and hardly vary across areas.

In line with H1a, we find that users express more negative sentiments towards immigrants on Twitter, when the saliency of immigration in the national news is higher in the week before (Table 2). When looking at the subsample of users who follow news on Twitter (Table 2, Model 3), we find support for H1b as well: higher saliency of immigration in the news outlets they follow is associated with more negativity towards immigrants in the week thereafter. A one standard deviation increase of the saliency of immigration in national news is associated with an increase in anti-immigrant sentiments of (0.86×0.023) 0.02. For personal news, the associated change is (3×0.0094) 0.03.

We further examined various time lags for news saliency (Table 3), using the same variables as included in Model 3. It appears that the impact of the saliency of immigration in national and personalized news outlets on sentiments users express on Twitter follows a common pattern. There is a strong association with news on immigration 1 day before, whereas news 3 days, 5 days, or 7 days ago matter much less. However, the *average* saliency of immigration in the week before, 2 weeks before, or 1 month before does have a statistically significant impact. This suggests that people respond to immediate news events (1-day lag) and to continuous news coverage of immigration over longer periods.

Returning to Table 2, we find that when Twitter users live in areas with a larger share of non-western

immigrants, they tend to tweet less negatively about immigrants ($b = -0.0062$, P less than 0.01, Model 1), which is in line with H2b, and against H2a. A one standard deviation increase of the non-western immigrant population is associated with a decrease in anti-immigrant sentiments in tweets of (10.63×0.0062) 0.066.

In line with H2c, we find that Twitter users who have a more-diverse Twitter network tend to tweet less negatively about immigrants. One standard deviation increase in the share of non-western immigrants that users follow is associated with a decline in negative sentiments of around 0.03 (Model 2) to 0.05 (Model 3). We further find no evidence for H3, which stated that the impact of the saliency of immigration in (national and personalized) media on sentiments is less strong in regions with a larger share of non-western immigrants.

Additional Analyses

We performed several robustness analyses. First, we run the models using SenticNet, which is an alternative sentiment algorithm (Cambria et al., 2016). Second, we analysed only negative sentiments with SentiStrength (excluding words expressing positive tone). Third, we re-examined regional effects, based on a small subgroup of Twitter users for which we were able to identify their geolocation with more precision. Overall, the results of these additional analyses confirm our conclusions drawn from the results presented in Table 2 (see Supplement for measurement and detailed results). The single exception to this is that in some of the models, the regional share of non-western immigrants is statistically insignificant.

As a further check, we estimated three more models to assess the consequences of potential biases (Table 4). One possible bias may arise when part of the sample belongs to ethnic minority groups, whereas the predictions are about ethnic majority members. Because there is no objective measure on the ethnicity of Twitter users, we relied on the method of ethnic identification of Twitter users described earlier. Model 4 presents the findings of the analysis which is based on the sub-sample of Twitter users which (presumably) belong to the ethnic majority population (47 per cent of total sample). This sub-sample excludes ethnic minority members (27 per cent) and those whose ethnicity could not be inferred (26 per cent). The findings confirm the earlier observations, except for the absence of an association between regional share of non-western immigrants and anti-immigrant sentiments.

Another source of bias could be the lack of appropriate control variables. We therefore added to our

Table 2. Multilevel hierarchical linear regression analysis of anti-immigrant sentiments on Twitter using SentiStrength (unstandardized coefficients)

	Model 1	Model 2	Model 3
<i>Fixed part</i>			
Intercept	-0.0904 ^{***} (0.0323)	-0.0740 ^{**} (0.0327)	0.0443 (0.0413)
Salience immigration in national news, 1 week	0.0281 ^{***} (0.0026)	0.0280 ^{***} (0.0026)	0.0230 ^{***} (0.0032)
Salience immigration in personalized news, 1 week			0.0094 ^{***} (0.0011)
Non-western immigrants in a region (per cent)	-0.0062 ^{**} (0.0027)	-0.0055 [*] (0.0028)	-0.0067 ^{**} (0.0030)
Salience immigration in national news, 1 week × non-western immigrants in a region (per cent)	0.0002 (0.0002)	0.0002 (0.0002)	0.0002 (0.0003)
Salience immigration in personalized news, 1 week × non-western immigrants in a region (per cent)			0.0000 (0.0001)
Non-western ties on Twitter (per cent)		-0.0016 ^{***} (0.0003)	-0.0025 ^{***} (0.0004)
<i>Sample size</i>			
N user-timepoints	314,589	314,589	223,685
N users	28,161	28,161	19,611
N regions	39	39	38
<i>Random part</i>			
σ_e^2	1.418250	1.418286	1.449103
σ_{u0}^2 (individual)	0.304811	0.303911	0.310972
σ_{v0}^2 (region)	0.002429	0.002604	0.002407
<i>Model fit</i>			
AIC	1,028,807	1,028,787	736,315
BIC	1,028,831	1,028,837	736,306

^{***} *P* less than 0.001.

^{**} *P* less than 0.01.

^{*} *P* less than 0.05 (two-tailed test).

Controls: population density region, number of followers on Twitter (log), number of connections on Twitter (log). Model 3 is based on the subsample of those who follow news outlets on Twitter.

main model two regional-level variables (per cent unemployed and per cent less than 35 years of age), and two individual-level variables: ethnicity and gender. We inferred gender of the Twitter users using the *genderComputer* algorithm (Vasilescu, Capiluppi and Serebrenik, 2014) (see Supplement for details). The findings, presented in Model 5, are in line with the results of models without these control variables. The evidence furthermore suggests that non-western ethnic minority members (as compared with western ethnic minority members or the ethnic majority) and women (as compared with men) express less-negative sentiments towards immigrants on Twitter.

It could also be that the personalized media exposure measure is flawed, because it does not incorporate the (personal) Twitter pages of journalists and reporters. We therefore constructed a new measure of personalized

media exposure to immigration, which is based on news from both official news outlets and journalists (see Supplement for measurement). Model 6 shows that the results remain the same: the more salient the topic of immigration is (in the week before) in the news pages that people follow on Twitter, the more negatively they tweet about immigration.

One key assumption made in this study is that when immigration becomes more salient in the news, people are aware of that, and as a result of increased salience and threat, they express more negative sentiments on Twitter. To further understand the linkage between the salience of immigration in the news and sentiments, we performed an analysis of whether people tweet (at day *t*) about the topic of immigration. If people are indeed aware of the increased salience of immigration in the news, one would expect them to tweet more often

Table 3. Multilevel hierarchical linear regression analysis of anti-immigrant attitudes on Twitter: testing different time lags and cumulative effects of national and personalized news on immigration

Specification	National news		Personalized news	
	b	S.E.	b	S.E.
1-day lag	0.0163***	0.0016	0.0027***	0.0005
3-days lag	0.0017	0.0016	0.0018	0.0005
5-days lag	0.0048***	0.0017	0.0003	0.0005
7-days lag	0.0048***	0.0016	0.0011	0.0006
1-week average	0.0272***	0.0026	0.0095***	0.0011
2-weeks average	0.0219***	0.0032	0.0132***	0.0015
1-month average	0.0401***	0.0044	0.0207***	0.0020

***P less than 0.001.

**P less than 0.01.

*P less than 0.05 (two-tailed test).

Controls: per cent non-western region, population density region, number of followers on Twitter (log), number of connections on Twitter (log).

about immigration. The results of this logistic regression model confirm this assumption: when immigration becomes more salient in the news, users increasingly tweet about the topic of immigration (Table 5). We find this association for both the national news environment (i.e. *The Guardian* and *The Sun*) and personalized news consumption.

Discussion and Conclusion

The growing ethnic diversity in Europe has resulted in concerns about anti-immigrant sentiments, discrimination, and racism. To date, most studies use survey data to capture people's prejudiced views. In a world that has become digitally connected, more research is needed on the views people publicly express about outgroups on social media platforms like Twitter, Facebook, and YouTube.

In this study, we argued that online data sources provide a rich tool for sociologists to monitor and study intergroup sentiments. We introduced and illustrated a new method of collecting data on online sentiments. We created a panel of 28,000 Twitter users in 39 regions in the United Kingdom. We applied automated text analysis to quantify anti-immigrant sentiments of 500,000 tweets over a 1-year period, and additionally studied the volume of tweets about immigration.

Results show that anti-immigrants strongly fluctuate over time. In support of group threat theory, and in line with earlier findings using surveys (Czymara and Dochow, 2018), our study provides evidence to suggest that people tend to tweet more negatively about

immigration in periods following more salient coverage of immigration in the news. We find this association both for national news coverage, and for the salience of immigration in the personalized set of outlets they follow on Twitter. Further analysis reveals an immediacy effect (1-day lag), and the impact of prolonged periods of news about immigration. We also find that stronger visibility of immigration in the news increases not only the negativity of tweets about immigration in the days and weeks following, but also the volume of tweets about immigration.

Other findings are in line with contact theory. We find some—though not conclusive—evidence to suggest that Twitter users living in areas with more non-western immigrants tend to tweet less negatively about immigrants. Looking more closely at people's connections on Twitter, we find that anti-immigrant sentiments are less common among users who follow a larger share of non-western immigrants.

We see several shortcomings of our study and opportunities for future research. First, this study did not consider the dynamics of the Twitter platform in detail, such as liking, retweeting, the non-randomness of the time line on Twitter, the posts of people users are connected to, to name only a few. Second, we did not compare our regional and longitudinal findings on Twitter with other sources, such as surveys or crime data. Recent work shows that anti-Black and anti-Muslim social media posts are predictive of offline racially and religiously aggravated crime (Williams et al., 2020). Studies in other fields report that opinion mining based on Twitter data strongly correlates with responses obtained via survey methods (Pasek et al., 2018).

Third, a disadvantage of Twitter is that the data are 'thin' and do not contain key individual-level data, such as gender, ethnicity, and education. However, in this study, we showed that, following recent developments in the field (Sloan et al., 2015; McCormick et al., 2017), one can apply sophisticated algorithms to infer demographic information about Twitter users. Our findings show that non-western ethnic minority members and women express less-negative sentiments towards immigrants on Twitter. These findings are reassuring, because they mirror findings from the survey literature, which equally reveal that anti-immigrant attitudes and prejudice are less strong among (non-western) ethnic minority members (Van der Zwan, Bles and Lubbers, 2017) and women (Ponce, 2017). Future research is encouraged to combine demographically-enriched Twitter data with methods of post-hoc stratification, to make the sample of Twitter users more representative of the general population.

Table 4. Multilevel hierarchical linear regression analysis of anti-immigrant sentiments on Twitter using SentiStrength (unstandardized coefficients): additional analyses

	Model 4 (ethnic majority)	Model 5 (control variables)	Model 6 (journalists)
<i>Fixed part</i>			
Intercept	0.0154 (0.0597)	0.1277** (0.0426)	0.1397 (0.0807)
Salience immigration in national news, 1 week	0.0278*** (0.0048)	0.0233*** (0.0032)	0.0241*** (0.0046)
Salience immigration in personalized news, 1 week	0.0102*** (0.0018)	0.0091*** (0.0011)	0.0111*** (0.0019)
Non-western immigrants in a region (per cent)	-0.0032 (0.0039)	-0.0064* (0.0032)	-0.0067 (0.0040)
Salience immigration in national news, 1 week × non-western immigrants in a region (per cent)	0.0002 (0.0005)	0.0002 (0.0003)	0.0002 (0.0004)
Salience immigration in personalized news, 1 week × non-western immigrants in a region (per cent)	0.0001 (0.0002)	-0.0000 (0.0001)	-0.0002 (0.0002)
Non-western ties on Twitter (per cent)	-0.0046*** (0.0008)	-0.0014*** (0.0004)	-0.0022** (0.0007)
Population density region	0.0000 (0.0000)	0.0000 (0.0000)	0.0000 (0.0000)
Unemployed in a region (per cent)		-0.0015 (0.0125)	-0.0112 (0.0150)
Age less than 35 in a region (per cent)		-0.0014 (0.0041)	0.0055 (0.0047)
Sex (ref=male)			
Female		-0.2155*** (0.0142)	-0.1876*** (0.0214)
Unknown		-0.1997*** (0.0133)	-0.1589*** (0.0207)
Ethnicity (ref=Western, including British)			
Non-Western ethnic minority		-0.0474*** (0.0138)	-0.0522* (0.0209)
Unknown		0.0869*** (0.0148)	0.0527* (0.0236)
<i>Sample size</i>			
N user-timepoints	102,049	221,196	103,363
N users	9764	19405	7891
N regions	38	38	38
<i>Random part</i>			
σ^2_c	1.456376	1.449408	1.473951
σ^2_{u0} (individual)	0.308703	0.301708	0.286947
σ^2_{v0} (region)	0.003182	0.002753	0.002939
<i>Model fit</i>			
AIC	336,692	727,720	341,013
BIC	336,816	727,916	341,214

***P less than 0.001.

**P less than 0.01.

*P less than 0.05 (two-tailed test).

Controls: number of followers on Twitter (log), number of connections on Twitter (log).

Overall, the results of this study underscore the argument that hypotheses derived from existing sociological theories work well for predicting online intergroup

sentiments. Group threat theory and contact theory can be used to develop a theoretical framework for understanding sentiments expressed online, despite the well-

Table 5. Multilevel logistic regression analysis of tweeting at day t about the topic of immigration on Twitter (unstandardized coefficients)

	Model 1	Model 2	Model 3
Fixed part			
Intercept	-4.9774*** (0.0451)	-5.1257*** (0.0440)	-5.1635*** (0.0561)
Salience immigration in national news, 1 week	0.0188*** (0.0023)	0.0188*** (0.0023)	0.0132*** (0.0027)
Salience immigration in personalized news, 1 week			0.0143*** (0.0009)
Non-western immigrants in region (per cent)	0.0106*** (0.0038)	0.0062* (0.0035)	0.0068* (0.0038)
Salience immigration in national news, 1 week × non-western immigrants in a region (per cent)	-0.0004** (0.0002)	-0.0004* (0.0002)	-0.0004* (0.0003)
Salience immigration in personalized news, one week × non-western immigrants in a region (per cent)			-0.0001 (0.0001)
Non-western ties on Twitter (per cent)		0.0116*** (0.0004)	0.0132*** (0.0005)
Sample size			
N user-timepoints	10,110,517	10,110,517	7,337,601
N users	28,161	28,161	20,439
N regions	39	39	38
Random part			
σ^2_{u0} (individual)	1.048767	1.010662	1.0154
σ^2_{v0} (region)	0.004095	0.003202	0.0028
Model fit			
AIC	2,369,078	2,368,165	1,753,141.6
BIC	2,369,205	2,368,307	1,753,307.3

*** P less than 0.001.** P less than 0.01.* P less than 0.05 (two-tailed test).

Controls: population density region, number of followers on Twitter (log), number of connections on Twitter (log). Model 3 is based on the subsample of those who follow news outlets on Twitter.

known ‘biases’ that arise when studying the online world: only a selective group is active on Twitter and publicly expressed sentiments may not correspond to private opinions. Our study finds patterns that resemble findings from the survey tradition on (private) outgroup attitudes. We believe that with the method illustrated in this study, Twitter and other social media platforms may become another tool for sociologists to study intergroup sentiments.

Notes

- 1 <https://www.abc.org.uk/data/>
- 2 It should be emphasized that in this study, we focus on the (time-varying) salience of immigration in the news Twitter users consume. We assume such time-varying changes in news salience is mostly exogenous. We do not examine the tone or framing of immigration

messages in news outlets (Van Klingeren et al., 2015) because of endogeneity issues that may arise. Specific news outlets people follow on Twitter may differ in their tone and framing of immigration, and people may select themselves to news outlets which are more in line with their prior beliefs.

Supplementary Data

Supplementary data are available at *ESR* online.

Acknowledgements

The authors are grateful to Ineke Maas for her valuable comments on earlier versions of this article. They would also like to thank Jonathan de Bruin and Parisa Zahedi for their assistance with data collection.

References

- Allport, G. W. (1954). *The Nature of Prejudice*. Reading, MA: Addison-Wesley.
- Alrababa'h, A. et al. (2021). Can exposure to celebrities reduce prejudice? The effect of Mohamed Salah on islamophobic behaviors and attitudes. *American Political Science Review*, 115, 1111–1128.
- Bakshy, E., Messing, S. and Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348, 1130–1132.
- Berning, C. C., Lubbers, M. and Schlueter, E. (2019). Media attention and radical right-wing populist party sympathy: longitudinal evidence from the Netherlands. *International Journal of Public Opinion Research*, 31, 93–120.
- Blalock, H. M. (1967). *Toward a Theory of Minority Group Relations*. New York, NY: Wiley.
- Blank, G. and Lutz, C. (2017). Representativeness of social media in Great Britain: investigating Facebook, LinkedIn, Twitter, Pinterest, Google+, and Instagram. *American Behavioral Scientist*, 61, 741–756.
- Bleich, E. et al. (2015). Media portrayals of minorities: Muslims in British newspaper headlines, 2001–2012. *Journal of Ethnic and Migration Studies*, 41, 942–962.
- Bliuc, A.-M. et al. (2018). Online networks of racial hate: a systematic review of 10 years of research on cyber-racism. *Computers in Human Behavior*, 87, 75–86.
- Blumer, H. (1958). Race prejudice as a sense of group position. *Pacific Sociological Review*, 1, 3–7.
- Boer, M. and van Tubergen, F. (2019). Media messages and attitudes toward Muslims and ethnic minorities: a panel study among ethnic majority adolescents in the Netherlands. *Social Science Research*, 83, 102311.
- Bohman, A. and Hjerm, M. (2016). In the wake of radical right electoral success: a cross-country comparative study of anti-immigration attitudes over time. *Journal of Ethnic and Migration Studies*, 42, 1729–1747.
- Boomgaarden, H. G. and Vliegenthart, R. (2009). How news content influences anti-immigration attitudes: Germany, 1993–2005. *European Journal of Political Research*, 48, 516–542.
- Cambria, E. et al. (2016). SenticNet 4: a semantic resource for sentiment analysis based on conceptual primitives. In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 2666–2677.
- Ceobanu, A. M. and Escandell, X. (2010). Comparative analyses of public attitudes toward immigrants and immigration using multinational survey data: a review of theories and research. *Annual Review of Sociology*, 36, 309–328.
- Cody, E. M. et al. (2015). Climate change sentiment on Twitter: an unsolicited public opinion poll. *PLoS One*, 10, e0136092.
- Czymara, C. S. and Dochow, S. (2018). Mass media and concerns about immigration in Germany in the 21st century: individual-level evidence over 15 years. *European Sociological Review*, 34, 381–401.
- Dovidio, J. F., Gaertner, S. L. and Kawakami, K. (2003). Intergroup contact: the past, present, and the future. *Group Processes & Intergroup Relations*, 6, 5–21.
- Dunaway, J. et al. (2011). Rebuilding or intruding? Media coverage and public opinion on Latino immigration in post-Katrina Louisiana. *Social Science Quarterly*, 92, 917–937.
- Eberl, J.-M. et al. (2018). The European media discourse on immigration and its effects: a literature review. *Annals of the International Communication Association*, 42, 207–223.
- Eurostat. (2018). Population density by NUTS 2 region (2008-2019) [Data file]. Retrieved December 1, 2020, from: <<https://ec.europa.eu/eurostat/databrowser/view/tgs00024/default/table?lang=en>>.
- Flores, R. D. (2017). Do anti-immigrant laws shape public sentiment? A study of Arizona's SB 1070 using Twitter data. *American Journal of Sociology*, 123, 333–384.
- Gorodzeisky, A. and Semyonov, M. (2016). Not only competitive threat but also racial prejudice: sources of anti-immigrant attitudes in European societies. *International Journal of Public Opinion Research*, 28, 331–354.
- Hainmueller, J. and Hiscox, M. J. (2010). Attitudes toward highly skilled and low-skilled immigration: evidence from a survey experiment. *American Political Science Review*, 104, 61–84.
- Hung, Z. (2020). SenticStrength, available from: <<https://github.com/zhunhung/Python-SentiStrength>>.
- Jacobs, L. and van der Linden, M. (2018). Tone matters: effects of exposure to positive and negative tone of television news stories on anti-immigrant attitudes and carry-over effects to uninvolved immigrant groups. *International Journal of Public Opinion Research*, 30, 211–232.
- Kearney, M. (2019). rtweet: Collecting and analyzing Twitter data. *Journal of Open Source Software*, 4, 1829–1829.
- Kuntz, A., Davidov, E. and Semyonov, M. (2017). The dynamic relations between economic conditions and anti-immigrant sentiment: a natural experiment in times of the European economic crisis. *International Journal of Comparative Sociology*, 58, 392–415.
- Kuran, T. (1997). *Private Truths, Public Lies*, Harvard University Press.
- Kuznetsova, A., Brockhoff, P. B. and Christensen, R. H. B. (2017). lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software*, 82, 1–26.
- McCormick, T. H. et al. (2017). Using Twitter for demographic and social science research: tools for data collection and processing. *Sociological Methods & Research*, 46, 390–421.
- McLaren, L., Boomgaarden, H. and Vliegenthart, R. (2018). News coverage and public concern about immigration in Britain. *International Journal of Public Opinion Research*, 30, 173–193.
- Mellon, J. and Prosser, C. (2017). Twitter and Facebook are not representative of the general population: political attitudes and demographics of British social media users. *Research & Politics*, 4, 2053168017720008.
- Meuleman, B. et al. (2020). Economic conditions, group relative deprivation and ethnic threat perceptions: a cross-national

- perspective. *Journal of Ethnic and Migration Studies*, 46, 593–611.
- Meuleman, B., Davidov, E. and Billiet, J. (2009). Changing attitudes toward immigration in Europe, 2002–2007: a dynamic group conflict theory approach. *Social Science Research*, 38, 352–365.
- Mottl, D. (2018). GetOldTweets3, available from: <<https://github.com/Mottl/GetOldTweets3>>.
- Office for National Statistics, Population of the UK by country of birth and nationality (2018). Population in the United Kingdom, excluding some residents in communal establishments, by country of birth [Data file]. Retrieved December 1, 2020, from <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/internationalmigration/datasets/populationoftheunitedkingdombycountryofbirthandnationality>
- Pasek, J. *et al.* (2018). The stability of economic correlations over time-identifying conditions under which survey tracking polls and Twitter sentiment yield similar conclusions. *Public Opinion Quarterly*, 82, 470–492.
- Pettigrew, T. and Tropp, L. (2008). How Does Intergroup Contact Reduce Prejudice? Meta-analytic Tests of Three Mediators. *European Journal of Social Psychology*, 38, 922–934.
- Ponce, A. (2017). Gender and anti-immigrant attitudes in Europe. *Socius*, 3, 1–17.
- Pottie-Sherman, Y. and Wilkes, R. (2017). Does size really matter? On the relationship between immigrant group size and anti-immigrant prejudice. *International Migration Review*, 51, 218–250.
- Ribeiro, F. N. *et al.* (2016). Sentibench-a benchmark comparison of state-of-the-practice sentiment analysis methods. *EPJ Data Science*, 5, 1–29.
- Savelkoul, M. *et al.* (2012). Comparing levels of anti-Muslim attitudes across Western countries. *Quality & Quantity*, 46, 1617–1624.
- Scheepers, P., Gijsberts, M. and Coenders, M. (2002). Ethnic exclusionism in European countries. Public opposition to civil rights for legal migrants as a response to perceived ethnic threat. *European Sociological Review*, 18, 17–34.
- Schemer, C. (2012). The influence of news media on stereotypic attitudes toward immigrants in a political campaign. *Journal of Communication*, 62, 739–757.
- Scheufele, D. A. and Tewksbury, D. (2007). Framing, agenda setting, and priming: the evolution of three media effects models. *Journal of Communication*, 57, 9–20.
- Schlueter, E. and Davidov, E. (2013). Contextual sources of perceived group threat: negative immigration-related news reports, immigrant group size and their interaction, Spain 1996–2007. *European Sociological Review*, 29, 179–191.
- Schlueter, E. and Scheepers, P. (2010). The relationship between outgroup size and anti-outgroup attitudes: a theoretical synthesis and empirical test of group threat- and intergroup contact theory. *Social Science Research*, 39, 285–295.
- Semyonov, M., Raijman, R. and Gorodzeisky, A. (2006). The rise of anti-foreigner sentiment in European societies (1988–2000). *American Sociological Review*, 71, 426–449.
- Sloan, L. *et al.* (2015). Who tweets? Deriving the demographic characteristics of age, occupation and social class from Twitter user meta-data. *PLoS One*, 10, e0115545.
- Sood, G. and Laohaprapanon, S. (2018). Predicting Race and Ethnicity From the Sequence of Characters in a Name. arXiv preprint arXiv:1805.02109.
- Thelwall, M. (2017). Sentiment analysis. In Sloan, L., and Quan-Haase, A. (Eds.), *The SAGE Handbook of Social Media Research Methods*. Thousand Oaks, CA: Sage, pp. 545–585.
- Van der Zwan, R., Bles, P. and Lubbers, M. (2017). Perceived migrant threat among migrants in Europe. *European Sociological Review*, 33, 518–533.
- Van Klingeren, M. *et al.* (2015). Real world is not enough: the media as an additional source of negative attitudes toward immigration, comparing Denmark and the Netherlands. *European Sociological Review*, 31, 268–283.
- Vasilescu, B., Capiluppi, A. and Serebrenik, A. (2014). Gender, representation and online participation: a quantitative study. *Interacting with Computers*, 26, 488–511.
- Weber, H. (2015). National and regional proportion of immigrants and perceived threat of immigration: a three-level analysis in Western Europe. *International Journal of Comparative Sociology*, 56, 116–140.
- Williams, M. (2019). *Hatred Behind the Screens: A Report on the Rise of Online Hate Speech*. Mishcon de Reya, available from: <<http://orca.cf.ac.uk/127085/>>.
- Williams, M. L. and Burnap, P. (2016). Cyberhate on social media in the aftermath of Woolwich: a case study in computational criminology and big data. *The British Journal of Criminology*, 56, 211–238.
- Williams, M. L. *et al.* (2020). Hate in the machine: anti-black and anti-Muslim social media posts as predictors of offline racially and religiously aggravated crime. *The British Journal of Criminology*, 60, 242–117.
- Zucker, H. G. (1978). The variable nature of news media influence. *Annals of the International Communication Association*, 2, 225–240.

Anastasia Menshikova is a PhD student at the Institute for Analytical Sociology, Linköping University, Sweden. Her research interests focus on biases and discriminatory attitudes in digital trace data and computational text analysis.

Frank van Tubergen is a professor of sociology at Utrecht University, Netherlands. His research interests include immigration, demography, religion and big data. He is a fellow of the European Academy of Sociology and the author of ‘Introduction to Sociology’ (Routledge, 2020).