



May 2018

No.373

**Fanning the Flames of Hate:  
Social Media and Hate Crime**

Karsten Müller and Carlo Schwarz

**WORKING PAPER SERIES**

Centre for Competitive Advantage in the Global Economy

Department of Economics



# Fanning the Flames of Hate: Social Media and Hate Crime\*

Karsten Müller<sup>†</sup> and Carlo Schwarz<sup>‡</sup>

May 21, 2018

## Abstract

This paper investigates the link between social media and hate crime using Facebook data. We study the case of Germany, where the recently emerged right-wing party *Alternative für Deutschland* (AfD) has developed a major social media presence. We show that right-wing anti-refugee sentiment on Facebook predicts violent crimes against refugees in otherwise similar municipalities with higher social media usage. To further establish causality, we exploit exogenous variation in major internet and Facebook outages, which fully undo the correlation between social media and hate crime. We further find that the effect decreases with distracting news events; increases with user network interactions; and does not hold for posts unrelated to refugees. Our results suggest that social media can act as a propagation mechanism between online hate speech and real-life violent crime.

JEL classification: D74, J15, Z10, D72, O35, N32, N34.

Keywords: social media, hate crime, minorities, Germany, AfD

---

\*We are grateful to Mirko Draca, Sascha Becker, Christopher Blattman, Leonardo Bursztyn, Evan Fradkin, Thiemo Fetzer, Magne Mogstad, Fabian Waldinger, and seminar participants at the University of Chicago and the Leverhulme Causality Conference at the University of Warwick for their helpful suggestions. We would also like to thank the Amadeu Antonio Stiftung for sharing their data on refugee attacks with us. Müller was supported by a Doctoral Training Centre scholarship granted by the Economic and Social Research Council [grant number 1500313]. Schwarz was supported by a Doctoral Scholarships from the Leverhulme Trust as part of the *Bridges* program.

<sup>†</sup>Warwick Business School, University of Warwick, k.muller@warwick.ac.uk

<sup>‡</sup>Department of Economics, University of Warwick, Centre for Competitive Advantage in the Global Economy (CAGE), c.r.schwarz@warwick.ac.uk

# 1 Introduction

Social media has come under increasing scrutiny during the most recent presidential election in the United States and all over the globe. Relatively recent phenomena such as fake news, social media echo chambers or bot farms have been subject of widespread media coverage and public discourse (e.g. New York Times, 2016, 2017a). The role of hate speech, especially online, has been at the center of a particularly intense and polarized debate. Despite the public interest and calls for policy action, there is little empirical evidence on the transmission of social media hate speech into real-life behavior.

In this paper, we study the link between social media and hate crime drawing on data from Facebook, the largest social media network. In particular, we investigate the relationship between anti-refugee sentiment on Facebook and hate crimes against refugees in Germany. The German setting is motivated by the relatively large recent influx of refugees into the country and the unfortunate frequency of violent crimes committed against them (see, for example, recent video coverage by New York Times, 2017b).

We create a measure for the salience of anti-refugee hate speech on social media based on the Facebook page of the “Alternative für Deutschland” (Alternative for Germany, AfD hereafter), a relatively new right-wing party that became the third-strongest faction in the German parliament following the 2017 federal election. The AfD has positioned itself as an anti-refugee and anti-immigration party; with more than 300,000 likes, it also has more followers than any other German party on Facebook (see Appendix A for a history of the AfD). This widespread reach makes the AfD’s Facebook page uniquely suited for an analysis of the relationship between social media and hate crime. In contrast to established political parties like Angela Merkel’s Christian Democratic Union (CDU) or the German Social Democrats (SPD), the AfD allows users to directly post messages on its Facebook wall; it is also the only party that does not explicitly outline rules of conduct, e.g. by threatening to remove racist, discriminating, or otherwise hateful comments. As a result, the AfD Facebook page contains far more posts and comments than those of other parties. With over 176,000 posts, more than 290,000 comments, and 500,000 likes by over 93,000 individual users, our data provide a unique insight into far-right Facebook usage and hate speech. This rich dataset enables us to understand the scale of the consumption of hate media online and analyze its real-world impact.

We posit that Germany-wide bursts of anti-refugee sentiment may push some potential perpetrators over the edge to carry out violent acts. If anti-refugee sentiment spreads through Facebook, we would expect hate crimes to be more likely to occur in municipalities with higher exposure to social

media, particularly when tensions are high. In our empirical strategy, we thus exploit differences in Facebook usage at the municipality level and weekly variation in posts about refugees on the AfD Facebook page – which we show are highly volatile – to study the effect of social media on hate crime. The inclusion of municipality and week fixed effects lets us abstract from average differences in anti-refugee sentiment across regions and time.

In our setting, the share of a municipality’s population that use the AfD Facebook page is an intuitive proxy for right-wing social media use; however, it is also correlated with differences in a host of observable municipality characteristics – most importantly the prevalence of right-wing ideology. We thus attempt to isolate the local component of social media usage that is uncorrelated with right-wing ideology by drawing on the number of users on the “Nutella Germany” page. With over 32 million likes, Nutella has one of the most popular Facebook pages in Germany and therefore provides a measure of general Facebook media use at the municipality level. While municipalities with high Nutella usage are more exposed to social media, they are not more likely to harbor right-wing attitudes.

Using these measures, we find that anti-refugee hate crimes increase disproportionately in areas with higher Facebook usage during periods of high anti-refugee sentiment online. This effect is especially pronounced for violent incidents against refugees, such as arson and assault. Taken at face value, this suggests a role for social media in the transmission of Germany-wide anti-refugee sentiment.

The main concern with this empirical strategy is that our measures of Facebook usage could be correlated with other observable and unobservable characteristics that explain the increase in anti-refugee hate crimes. We address these concerns twofold. First, the Nutella data also enable us to create a dummy for municipalities with many Nutella users *within a county*. Crucially, this measure of social media activity is balanced across a plethora of observable characteristics, most importantly general internet usage, voting patterns, education, “pull-factors” such as immigration and religious composition, and proxies of xenophobic attitudes. As such, our empirical strategy is related to Enikolopov et al. (2016), who use an instrument to generate variation in social media usage unrelated to observables. We further control for increases in anti-refugee incidents in weeks of high refugee salience that are driven by observable municipality differences. As it turns out, using our regression framework to “match” municipalities on these characteristics (Angrist and Pischke, 2008) makes little difference for our point estimates.

Second, we provide quasi-experimental evidence using local internet disruptions and Germany-wide Facebook outages to narrow down the social media transmission channel and rule out the influence

of potential unobservable factors. We draw on user-reported data to identify municipalities and weeks with large-scale internet outages, which we verify using news articles in regional and national media outlets. Notably, internet disruptions are geographically dispersed and orthogonal to AfD likes on Facebook, as well as measures of general internet affinity. This suggests that internet outages create exogenous variation in social media access *independent* of their impact through other online channels. Intuitively, we find that such outages decrease local social media activity in a sub-sample of geo-matched Facebook posts, comments, and likes. While higher anti-refugee sentiment increases the number of hate crimes, the effect is strongly reduced for municipalities experiencing internet outages in a given week. Quantitatively, we find that a typical internet disruption *fully* mediates the effect of social media salience on hate crime. Internet problems themselves do not have an independent effect on attacks, which makes it unlikely that we are capturing a “time use effect” of potential perpetrators fixing their internet access; this further points to social media as the propagation mechanism.

We also draw on a number of large, Germany-wide Facebook outages, which disrupted the ability of page owners to post content or users to interact with each other. These outages are not the result of internet problems but rather programming mistakes or server problems at Facebook. Consistent with a causal effect of social media, we again find that the effect of refugee posts on hate crimes essentially vanishes in weeks of major Facebook outages.

We further probe the social media transmission channel in additional tests. The effect we uncover appears to be limited to refugee-specific sentiment: other posts on the AfD Facebook page, e.g. those related to the European Union or Jews, do not have significant predictive power for hate crimes. Consistent with the hypothesis that hate speech is transmitted through social networks, the size of the effect on hate crime is higher in regions where AfD users show higher engagement on Facebook via likes and comments. Importantly, these engagement proxies are essentially orthogonal to social media usage and thus provide meaningful additional variation. In addition, we analyze how other salient news events mediate the effect of anti-refugee Facebook posts on the number of violent incidents, in the spirit of Eisensee and Strmberg (2007); Durante and Zhuravskaya (2016). Specifically, we look at the European Soccer Championship, Brexit, and Donald Trump’s presidential election, all of which attracted considerable attention in the German media. We find a significantly reduced effect in weeks in which these events partially crowd out the salience of refugees.

Moreover, we provide additional evidence using alternative proxies for refugee salience constructed from newspaper reports and right-wing protest participation. While our measure of social media sentiment is necessarily correlated with media attention about refugees, we find that coverage

by major German news outlets does not have an independent effect on hate crimes, once we take social media into account. General right-wing sentiment, as measured by the weekly number of protesters of the “Patriotic Europeans Against the Islamisation of the West” (PEGIDA), in turn particularly spurs *local* demonstrations, in contrast to the AfD-based salience that only has an effect on violent incidents. Because our preferred measures of social media exposure are orthogonal to internet usage and local right-wing ideology – and *negatively* correlated with newspaper consumption – these findings are consistent with an independent propagation effect of social media on hate crimes.

When interpreting our results, we do not claim that social media itself *causes* crimes against refugees out of thin air. In fact, hate crimes are likely to have many fundamental drivers; local differences in xenophobic ideology or a higher salience of immigrants are only two obvious examples. Rather, our argument is that social media can act as a propagating mechanism for the flare-up of hateful sentiments. Taken together, the evidence we present suggests that quasi-random shifts in the local population’s exposure to such sentiments on social media can magnify their effect on refugee attacks.

*Related literature.* Our work extends existing work in several ways. First, we build on the literature on media exposure and persuasion on real-life outcomes, particularly in regards to violence (see e.g. DellaVigna and Gentzkow, 2010; DellaVigna and La Ferrara, 2015). Most closely related is work by, among others, Yanagizawa-Drott (2014), Adena et al. (2015), and DellaVigna et al. (2014), who show how traditional media can trigger violent outbursts or ethnic hatred. Dahl and DellaVigna (2009) show that while exposure to violent movies can increase the propensity to be violent in experimental settings, it decreases violent crime in the field. Television has also been shown to be associated with short-lived outbursts of domestic violence (Card and Dahl, 2011). In other research, Bhuller et al. (2013) show that exposure to pornographic material on the internet is associated with increased sex crime.

We extend this literature by investigating the role of social media, which differs from traditional media in allowing user exchange and self-select into preferred topics and viewpoints. Sunstein (2009, 2017) argues that social media tends to limit the spectrum of information a person absorbs by creating echo chambers which reinforce similar ideas. Our findings provide evidence for real-life negative effects of social media filter bubbles. In line with Sunstein (2002), we find that the deliberation of extreme viewpoints on social media increases polarization. Previous research already documented the high prevalence of hate speech on social media (Oksanen et al., 2014) as well as the possibility to measure the racial animus in the US using Google search data (Stephens-Davidowitz, 2014). We show that hateful sentiments are not only propagated through these networks, but also associated with

more hate crimes. In other related research, Enikolopov et al. (2016) show that social media exposure spurs protest participation in Russia by reducing coordination costs and Gavazza et al. (2015) find that broadband diffusion decreased voter turnout in the United Kingdom (see also Gentzkow, 2006; Manacorda et al., 2017).

Second, the paper speaks to the literature on the polarization of citizens (e.g Fiorina and Abrams, 2008). Most of the existing literature thus far has failed to find empirical evidence for polarizing effects of social media (Boxell et al., 2017) or found that social media reduces polarization (Barberá, 2014). Our work implies that even if overall polarization is unaffected by social media, hateful content in online networks is associated with violent crimes.

Third, we contribute to the literature on culture and violence. Summarizing the vast area of research into cultural and religious fragmentation, Alesina and La Ferrara (2005) find that they predict the likelihood of civil war across countries. Voigtlander and Voth (2012) show that anti-semitic violence in Germany is highly persistent: pogroms during the era of the Black Death predict pogroms in the 1920s, Jewish deportations, and synagogue attacks during the rise of the Nazi party. Similarly, Jha (2013) shows that medieval interethnic complementarities in trade decrease the likelihood of modern Hindu-Muslim riots. These papers, however, are largely silent on the existence and effects of volatile, short-lived bursts in sentiment leading to violent incidents. As such, our work is particularly related to Fouka and Voth (2013), who show that monthly variation in public acrimony between Greek and German politicians during the Greek debt crisis affected German car purchases particularly in areas of Greece where German troops committed war crimes during World War II. Our results also fall in line with the findings of Colussi et al. (2016), who show that a higher salience of minority groups increases the likelihood of hate crimes.

While traditional media such as television are regulated in most countries, social media has only now moved into the focus of legislators. Our work is thus particularly topical in light of the political discussions in many countries about anti-hate speech laws and censoring of hate speech on social media. The German parliament, for example, passed an anti online hate speech law (“Netzwerkdurchsuchungsgesetz”) on June 30, 2017, which threatens to fine providers of online platforms such as Facebook with up to EUR 50 million for failing to delete “criminal” contents that are “obviously unlawful”. The controversial law passed parliament on the initiative of German Minister of Justice Heiko Maas, who lamented the unwillingness of social media platforms to address “online hate crime”.<sup>1</sup> The European Union has issued independent guidelines calling on social media companies to remove illegal hate

---

<sup>1</sup>See, for example, the official statement of the German parliament on bundestag.de.

speech as well. In the United Kingdom, the Crown Prosecution Service plans to increase prosecution of online hate crimes (The Guardian, 2017; BBC, 2017). Our paper serves as a first attempt to address this important topic empirically.

The paper proceeds as follows. In Section 2 we introduce the data used in our empirical analysis. Section 3 introduces basic correlations in the data as well as the empirical strategy and main empirical results. Section 4 concludes.

## **2 Data**

We construct a new data set of social media activity and anti-refugee hate crimes in Germany, centered around the country's most popular social media network: Facebook. In total, we combine data from 14 different sources which we describe in more detail in the following subsections: (1) Municipality-level data on anti-refugee hate crimes; (2) Facebook data on posts, likes, and comments from the AfD and Nutella pages; (3) hand-collected municipality-level data on Facebook user locations; (4) municipality-level data on internet outages; (5) a hand-coded dataset on major weekly Facebook outages; (6) socioeconomic data on the municipality and county level from the German Statistical Office; (7) election district voting data; (8) county-level data on broadband access; (9) survey data on internet usage from Eurostat; (10) municipality-level data on newspaper sales; (11) city-level data on neo-Nazi murders and historical anti-Semitism; (12) hand-collected data on refugee salience in traditional and right-wing media outlets; (13) weekly data on protest participation at PEGIDA; and (14) weekly Google search data on major news events in our sample. The final panel dataset covers 4,466 German municipalities for the 111 weeks from 1st January 2015 to 13th February 2017. Summary statistics for the main variables of interest can be found in Table 1. The online appendix provides a comprehensive overview of the data sources and variable definitions (see Table A.6) and summary statistics for additional control variables (see Table A.5).

### **2.1 Anti-Refugee Incidents**

The data on the number of incidents targeting refugees were collected by the Amadeu Antonio Foundation and Pro Asyl (a pro asylum NGO). They cover anti-refugee incidents in Germany ranging from anti-refugee graffiti, arson of refugee homes to assault. The period of January 2015 through early 2017 is of particular interest since it covers the beginning and height of the refugee crisis in Germany. All of the 3,335 incidents against refugees feature a short description and are classified into 5 groups.

**Table 1: SUMMARY STATISTICS FOR MAIN VARIABLES**


	Level	Obs.	Mean	S.D.	Min.	Max.
<b>Additional Media and Internet Controls<sup>†</sup></b>						
Refugee attacks/Refugees	Muni.-Week	495,726	0.038	1.006	0	227.273
Arson attacks/Refugees	Muni.-Week	495,726	0.002	0.125	0	19.960
Other property attack/Refugees	Muni.-Week	495,726	0.026	0.874	0	227.273
Assaults/Refugees	Muni.-Week	495,726	0.006	0.380	0	125.000
Demonstrations/Refugees	Muni.-Week	495,726	0.004	0.257	0	125.000
Suspected cases/Refugees	Muni.-Week	495,726	0.000	0.031	0	18.315
<b>Social Media Data</b>						
AfD users/Pop. <sup>†</sup>	Municipality	495,726	3.006	2.863	0	80.25223
Nutella users/Pop. <sup>†</sup>	Municipality	495,726	2.007	2.040	0	35.82689
<i>I<sub>Many Nutella Users</sub></i>	Municipality	495,726	0.417	0.493	0	1
Refugee posts	Week	495,726	84.027	61.422	2	259
Posts/AfD users	Municipality	395,493	0.554	3.882	0	118
Comments/AfD users	Municipality	395,493	1.085	7.251	0	270
Likes/AfD users	Municipality	395,493	1.760	12.263	0	370
<b>Auxiliary Variables</b>						
<i>I<sub>Internet outage</sub></i>	Muni.-Week	495,726	0.001	0.025	0	1
<i>I<sub>Facebook outage</sub></i>	Week	495,726	0.081	0.273	0	1
PEGIDA demonstrators	Week	393,008	4530	3931	1075	25000
Refugee news coverage	Week	495,726	0.000	1.917	-2.825	5.738
Right-wing news coverage	Week	495,726	3.982	2.959	0	11
<b>Baseline Controls</b>						
Population (2015) <sup>†</sup>	Municipality	495,726	1.840	7.478	0.034	352.003
GDP/Worker	County	493,617	63095	9846	46835	136763
Population density	Municipality	495,726	281.921	381.634	6.555	4653.184
AfD vote share (2017)	Election Distr.	495,726	14.216	5.987	4.915	35.019
Share Abitur	Municipality	495,726	29.038	8.251	0	58.466
Share broadband access	Municipality	495,726	82.999	10.656	43.500	100.000
Share immigrants	Municipality	483,072	13.962	7.627	1.819	49.722
<b>Raw Data</b>						
Refugee attacks	Muni.-Week	495,726	0.007	0.099	0	8
Refugees (2015) <sup>†</sup>	Municipality	495,726	0.230	0.201	0.004	4.965
AfD users	Municipality	495,726	7.700	49.881	0	2559
Nutella users	Municipality	495,726	4.915	27.005	0	1286

*Notes:* This table reports summary statistics for the variables in the estimation sample. Variables tagged with a † are scaled by population in 10,000. Share variables are in percent.

The most common incidents are property damage to refugee homes (2,226 incidents), followed by assault (534), anti-refugee demonstrations (339), arson (225). 11 incidents are classified as suspected cases which were still under investigation.

All incidents are geo-coded with an exact longitude and latitude. We use these coordinates to assign each incidents to a municipality.<sup>2</sup> Figure 1 shows the total number of anti-refugee incidents per asylum seekers in our observation period for each German municipality. The data appear to be of high quality. Each entry includes the source from which an incident was identified, usually national or local media outlets. We hand-checked a random sample of 100 incidents and found that the coding accurately reflected the information reported in the respective source.<sup>3</sup>

**Figure 1: ANTI-REFUGEE INCIDENTS PER ASYLUM SEEKERS, BY MUNICIPALITY**



*Notes:* This figure shows the number of anti-refugee incidents scaled by the number of asylum seekers (in 10,000) for each of the 4,466 German Municipalities in the sample. The refugee attack data were collected by the Amadeu Antonio foundation and PRO ASYL.


---

<sup>2</sup>For the assignment of coordinates to municipalities, we use the shape files provided by ©GeoBasis-DE/BKG 2016 website. Overall the shape file contains data for the 4,679 German municipalities (“Gemeindeverwaltungsverband”). 213 of these municipalities do not have any inhabitants (e.g. forest areas) nor anti-refugee incidents; hence, we only keep the remaining 4,466 municipalities in our estimation sample. We use the level of the “Gemeindeverwaltungsverband” since these exhibit smaller differences in their size and population in comparison to the 11,165 German “Gemeinden”, which makes it more suitable for spatial analysis according to the data provider (see link).

<sup>3</sup>Note that reporting bias is unlikely to play a role because our empirical estimations include municipality and week fixed effects. Any bias in reporting would thus have to vary by municipality *and* week as a function of social media and anti-refugeesentiment. This is particularly unlikely because, as we will show later, some of our measures of local social media penetration are unrelated to almost all municipality characteristics.

**Figure 2: AfD AND NUTELLA FACEBOOK USAGE PER CAPITA, BY MUNICIPALITY**

(a) AfD Facebook User



(b) Nutella Facebook User



*Notes:* These maps plot the number of Facebook users per capita (in 10,000) for each of the 4,466 German Municipalities as measured by the geo-located user data obtained from the Facebook pages of the Alternative for Germany and Nutella Germany.

## 2.2 Facebook Data on Refugee Salience

Next, we construct a measure for the salience of anti-refugee hate speech on social media. We start by using the Facebook Graph API to collect all status posts, comments and likes from the Facebook pages of the AfD (see Appendix B.1. for an introduction to Facebook). The Facebook Graph API provides a unique identifier for each post which allows us to assign comments and likes to posts. Additionally, the graph API provides unique identifiers for each user who either posts, comments, or likes anything on the page. Overall, we collected 176,153 posts, 290,854 comments, 510,268 likes, and 93,806 individual user IDs from the AfD Facebook page.

We use the AfD’s Facebook page because the AfD is by far the largest far-right party in Germany and has the highest number of Facebook followers of *any* of the German parties. This makes the AfD the arguably most important platform of exchange about refugees for Germany’s right-wing social media users. As our baseline measure for the salience of anti-refugee hate speech on social media, we use the number of posts on the AfD Facebook page in any given week that contain the word “Flüchtling” (refugee). We also construct analog measures based on comments or likes.

A potential downside of this procedure is that we may also tag posts that do not express negative

sentiments towards refugees. However, a careful reading of posts and comments reveals that the overwhelming majority appear to be in agreement with the positions of the AfD. This is perhaps unsurprising given that only people who “like” the AfD page on Facebook will be informed about new posts: critics, on the other hand, have a strong incentive not to indicate publicly that they “like” the party. Nevertheless, our measure is closely related to the general salience of refugees, which we explore in section 3.6.

To get an idea of the tone of exchange on the page, consider this example post: “Maybe there is a plot to exterminate the German genes with the large streams of refugees. But what could be the reason, [revenge for] WW2?”.<sup>4</sup> In Table A.1 in the online appendix, we show further representative examples of posts published on the AfD page.

We also construct measures for the salience of other topics by tagging posts containing the words “Islam”, “Muslim”, “Jude” (Jew) or “EU” on the AfD’s Facebook page.

### **2.3 Municipality-Level Facebook Usage Data**

For our empirical strategy, we construct a measure of social media usage at the municipality-level. Because survey data on Facebook usage, to our knowledge, are only available at the level of the 16 federal states, we hand-collect user location data by using the unique user identifiers provided by the Facebook Graph API. Due to Facebook’s privacy policy, we are only able to collect this information for people who make it publicly available.

Our first measure of Facebook usage is based on the users of the AfD Facebook page. In total we have 93,806 user ids.<sup>5</sup> We were able to hand-collect a place of residence for 39,632 users. Overall, we were able to identify at least one AfD Facebook user for 3,565 of the 4,466 municipalities.

Using the location data for AfD users, we can also assign posts, comments, likes to municipalities. Based on these data we construct auxiliary measures of social media reach, e.g. the number of local posts scaled over the number of AfD users. We find that some users post and comment excessively, which leads to a few outliers in measuring how active users are in a given municipality; we thus winsorize the number of posts, comments, and likes we can attribute to local users at the 99.9th percentile to avoid individual users driving the results.

---

<sup>4</sup>Original Post in German: “Evtl. soll ja die deutsche Genetik ausgerottet werden, durch große Flüchtlingsströme. Doch was könnte der Grund sein, WW2?”

<sup>5</sup>We can only collect data for users who have been active on the AfD Facebook page at least once. As a result, the total number of user ids is not identical to the more than 300,000 people who have liked the AfD Facebook page.

Next, we create a measure of general Facebook usage in Germany based on the page of Nutella Germany. The Nutella Germany page was chosen because, with more than 32 million followers, it is the second most popular Facebook page in Germany and should therefore provide an accurate measure of Facebook activity across municipalities.<sup>6</sup> We were able to collect 12,762 posts, which in turn received 38,002 comments and 51,465 likes; these reflect the actions of 63,207 individual users on the Nutella page. Using the same procedure described above, we hand-collect the place of residence for 21,915 users. Note that the number of user data for Nutella’s page is considerably lower despite the much higher number of Facebook fans because we can only collect data for users who at least posted, commented or liked something at least once. Nevertheless, we consider the almost 22,000 places of residence we collected as a good approximation of Facebook usage in Germany and a considerable improvement over existing survey measures. We have at least one Nutella user for 3,190 municipalities.

While the AfD and Nutella measures capture slightly distinct concepts, Figure A.3 in the online appendix shows that they are highly correlated.<sup>7</sup> In Figure A.3b, we show that this connection also holds for the unscaled number of users and at all points of the variable distributions. Figure 2 visualizes the number of AfD and Nutella Facebook users per capita for each German municipality.

As a third measure of social media usage, we create a dummy variable equal to 1 for municipalities that are in the top tercile of Nutella users per capita *within a county* ( $I_{Many\ Nutella\ Users}$ ). This measure has the advantage of being uncorrelated with a host of observable municipality characteristics, which we explore in more detail in Section 3.2. The dummy is again highly correlated with the proxy based on AfD users: the F-statistic of an OLS regression is above 95 (see Table A.9 in the online appendix).

## 2.4 Data on Internet and Facebook Outages

We collect data on local internet outages from Heise Online. Heise lists user reports of internet problems by dialing code areas, as well as their start time and the duration. We use the dialing codes to assign internet problems to municipalities; the start date and duration allow us to count the number

---

<sup>6</sup>The most popular Facebook page is the fan page of the football team FC Bayern Munich. We decided against collecting data from their page because we wanted to avoid users being locally concentrated in Bavaria.

<sup>7</sup>A simple OLS regression yields a highly significant coefficient with F-statistics of over 15 (see Table A.9 in the online appendix). Nutella also explains a substantial fraction of around 11% in the total variation of per-capita AfD users on Facebook.

of problems in each municipality in each week.<sup>8</sup> The reports of internet outages are geographically dispersed with no clear patterns of regional clustering (see Figure A.4).

To validate the Heise data, we search for newspaper reports on major internet disruptions. For all major disruptions, we identify in newspapers, the data suggests an increase in the number of outages in the same week which is specific to the provider encountering the disruption. Table A.3 lists several examples of news paper reports on internet outages and the respective information in our data.<sup>9</sup>

Because some reports may reflect individual users' glitches rather than general disruptions, we exclude reported outages with a duration of less than 24 hours.<sup>10</sup> Because we want to measure internet outages that affect a significant part of the population, we construct a dummy variable equal to 1 for municipality-weeks for the top quartile in the ratio of outages per capita. As we will discuss later, our results are robust to using alternative definitions of this cut-off.

In addition to the data on internet outages, we collect information on major Facebook disruptions. To identify these, we start by searching for newspaper reports of Facebook problems in our sample period. In total, we find reports on 8 large outages (see Table A.4 for an overview and more details). As an independent source of validation, we also obtain the number of weekly user-reported problems with Facebook on the Facebook page of "Allestörungen", a portal for aggregating user complaints on individual websites and apps. Perhaps unsurprisingly, the 8 outages widely reported on in the news media are also associated with spikes in user-reported problems.

Using these data, we define a dummy variable that is 1 for weeks with Facebook outages. These outages have the advantage that they are specific to Facebook; in fact, they are uncorrelated with the total number of weekly internet outages from our Heise data. The downside is that they are rare and, in contrast to the internet disruptions, only vary by week.

## 2.5 Auxiliary and Control Variables

We obtain control variables from a host of sources, which are explained in more detail in the online appendix. Socioeconomic data on the municipality and county level are from the German Statistical Office, available via [www.regionalstatistik.de](http://www.regionalstatistik.de). We include information on the population of each

---

<sup>8</sup>In the case where the dialing code areas span multiple municipalities, we assign an internet outage to the municipality that has the largest overlap with the dialing code area. This is preferable to assigning the outage to all municipalities the dialing code intersects with because some dialing code areas have minor overlaps with a large number of municipalities. Assigning the internet outage to all of these municipalities would introduce substantial noise.

<sup>9</sup>To interpret the number of outages, note that the Heise data reports an average of 4 reported internet outages per provider per week; hence, even an increase of 15 reported outages represents a large increase.

<sup>10</sup>In some cases, users do not seem to report the end date of the internet outage, which can lead to unlikely durations of several months. We thus winsorize the maximum duration at 3 weeks, but this choice is not material for our results.

municipality split by age group, GDP per worker, population density, the share of the population with a high school degree (“Abitur”), the share of the population receiving social benefits, and the share working in manufacturing. To control for “pull factors” of anti-minority crimes, we also obtain the share of the population that are immigrants and the number of asylum seekers; we use the latter to scale the number of refugee attacks in our main specification. We collect vote results data for the 2017 German Federal Election at the election district level from [www.bundeswahlleiter.de](http://www.bundeswahlleiter.de), which contain data on vote shares and voter turnout.

To measure the extent to which people use the internet in different localities, we use the share of households in a county with broadband access, collected by the Federal Ministry of Transport and Digital Infrastructure. Broadband access is highly correlated with publicly available survey data on individuals’ internet use from Eurostat, which is only available on the state-level (see Figure A.5 in the online appendix).<sup>11</sup> This suggests that broadband access is a sound proxy for local differences in internet usage. In addition, we use the number of registered .de internet domains per capita in a county to measure internet affinity, which has a correlation of around 0.39 with broadband access.

To measure the local penetration of traditional media, we obtain data for newspaper sales in municipalities for the years 2016/2017 from the “Zeitungsmarktforschung Gesellschaft der deutschen Zeitungen (ZMG)” (Society for Market Research of German Newspapers). These data contain the number of print newspapers sold in each municipality with more than 3,000 inhabitants. Newspapers are listed if, in any given town, they sell (1) at least 50 copies and (2) have a market share of at least 1%. Based on this data, we construct a measure of traditional newspaper consumption as the number of newspaper sales per capita.

As a measure for the local prevalence of right-wing extremism we use the number of murders committed by neo-Nazis in each municipality from 1990 until 2016. These data were collected by the project “Mut gegen rechte Gewalt” (Courage Against Right-Wing Violence). We complement this proxy for contemporary right-wing violence with data on the historic prevalence of anti-semitism collected by Voigtlander and Voth (2012). From their dataset, we use the natural logarithm of one plus the number of deported Jews as well as one plus the number of letters written to “Der Stürmer”, the antisemitic newspaper published by Nazi politician Julius Streicher.<sup>12</sup>

---

<sup>11</sup>In particular, we use the share of households with access to internet speeds of 16 Mbit/s and above, which is in the middle of the five available maximum speed categories from the Ministry’s data (above 2, 6, 16, 30, 50 Mbit/s). We focus on the 16 Mbit/s cut-off because it has the highest correlation with actual internet use data; using the other measures instead has no bearing on the results.

<sup>12</sup>Note that we use the log-numbers instead of scaled variables as controls because the data from Voigtlander and Voth (2012) only cover a fraction of the municipalities in our sample. We code cities with no information on deported Jews and

We further construct several alternative proxies for refugee salience in Germany. First, we collect the number of news items containing the word “Flüchtling” (refugee) from the websites of major newspapers in Germany, namely the “Frankfurter Allgemeine Zeitung”, “Spiegel”, “Zeit”, and “Handelsblatt”.<sup>13</sup> For each of these news sites we count the number of news items about refugees in each week. We repeat the same procedure for the “Tagesschau” website. The “Tagesschau” is the most important news show on German television and further maintains a highly-frequented website; the number of news items reported there thus provide a proxy for refugee salience on television. We take the first principal component of these news measures as the baseline proxy for media coverage on refugees.<sup>14</sup> Finally, we include news items of the far-right newspaper “Compact Magazin” as a measure of refugee salience in right-wing news outside of social media.<sup>15</sup>

As a further measure of right-wing sentiment in Germany that is not based on media mentions, we obtain the numbers of weekly PEGIDA demonstrators (Patriotic Europeans against the Islamization of the West) from the blog and Twitter account [durchgezaehlt.org](http://durchgezaehlt.org), a research group focused on crowd size estimation based at the University of Leipzig. The PEGIDA demonstrations took place on a weekly basis every Monday in Dresden from October 2014 until 2017; we do not have data for five dates in late 2016 and early 2017 and treat these as missing.

Finally, we obtain data from Google trends for the overall interest in the search terms “Brexit”, “Trump” and “UEFA Euro 2016” in Germany to proxy for distracting news events. Google scales the weekly number of searches for these terms on a scale from 0 to 100, where 100 marks the week with the highest search interest in the preceding 5 years. Time series plots suggest that these measures are sound approximations for attention paid to Brexit, the Trump election, and the UEFA European Championship (one of the most widely followed sports events in Germany).

---

Stürmer letters as zero.

<sup>13</sup>We cover most of the widely-read nationwide newspapers based on circulation data published by the ZMG on their website. We also attempted to collect articles from the websites of “Bild”, “Süddeutsche Zeitung” and “Welt”. Unfortunately the search results for their websites do not go back far enough to cover our observation period.

<sup>14</sup>In the online appendix, we show that our results using this measure are also robust to using the individual newspapers’ article numbers.

<sup>15</sup>Compact was at the center of a controversy for hosting a booth at the Leipzig Book Fair in March 2016, which led to a public outcry and protests, both of which received considerable media attention in Germany.

## 3 The Effects of Social Media on Hate Crime: Evidence from the German Far Right

### 3.1 Introductory Correlations

We start our analysis by documenting simple correlations between social media and attacks on refugees in Germany. The results in this section should thus be interpreted as purely suggestive and do not allow for causal inference; nevertheless, we consider the findings insightful, because we are not aware of previous empirical evidence on the cross-sectional and time series relationships between social media and hate crime.


First, we plot the total number of posts on refugees on the AfD's Facebook page against the number of anti-refugee incidents in Figure 3a. Weeks with more anti-refugee posts also tend to have more anti-refugee incidents. The correlation also holds in a time series regression of refugee attacks on AfD posts, which yields an adjusted  $R^2$  of around 34% (reported in online appendix Table A.8). The next relevant question is whether refugee attacks are concentrated in areas with more right-wing social media users. Figure 3b shows the share of municipalities with at least a single refugee attack, depending on whether we can identify *at least one* AfD user. The share of municipalities with attacks is around 10% for municipalities without and around one third for municipalities with AfD users. Out of the total 3,335 attacks on refugees in our sample, 3,171 occurred in municipalities with AfD users. A  $t$ -test rejects the null hypothesis of no difference between the mean of the two groups with a value of 5.292.

To put the time series and cross-sectional perspective together, we re-estimate the time series regressions separately for municipalities with at least 1 user on the AfD page and contrast it with those where we have zero users. To make the coefficients comparable, we standardize the number of refugee attacks in each sample to have a mean of 0 and standard deviation of 1. If online transmission matters, Germany-wide sentiment should have a larger propagating effect on violence in municipalities with AfD users. The regression results show that the time series relationship is much stronger where AfD users are present, yielding an almost 80% higher  $R^2$  (see Table A.8 in the online appendix). The (standardized) coefficient in the sample with at least one AfD user is around a third larger than that for the zero-user sample.

The correlations documented in this section are suggestive of a strong link between anti-refugee posts on social media and hate crimes. However, it is difficult to draw conclusions from these results. In particular, it remains unclear how much of the effect is driven by users on Facebook *reacting* to,


### Figure 3: Introductory Correlations

(a) Anti-Refugee Posts and Incidents over time



Notes: This figure shows the number of anti-refugee post on the Facebook page of the “Alternative for Germany” and the number of anti-refugee incidents in Germany over time.

(b) Share of Municipalities with Refugee Attacks, by AfD Users



Notes: This figure plots the share of municipalities with at least one refugee attack in our sample by whether we have evidence of at least one AfD user in the municipality. For 3,563 municipalities we are able to identify at least one AfD user. For 903 municipalities we find no AfD user.

rather than causing violence. It is also unclear to which extent common shocks such as the news cycle or persistent local attitudes towards minorities affect both the online and offline behavior we are observing. We next push these initial findings further by developing an empirical strategy that addresses concerns about reverse causality and omitted time and municipality factors.

### 3.2 Empirical Strategy

To investigate the effect of social media on anti-refugee incidents, we run fixed effects panel regressions akin to a Bartik-type approach (Bartik, 1991). In particular, we use the interaction of Facebook users per capita in municipality  $m$  to measure exposure to the Germany-wide salience of anti-refugee hate speech in week  $w$  ( $Refugee Posts_w$ ), which we collect from the AfD’s Facebook page. This empirical set-up creates variation by week and municipality, which we exploit in the following regression model:

$$Refugee Attacks/Refugees_{mw} = \beta Refugee Posts_w \times Social Media Users/Pop._m + \gamma Controls_m \times Refugee Posts_w + Week FE_w + Municipality FE_m + \epsilon_{mw}, \quad (1)$$

In this specification,  $\beta$  measures the differential effect of one additional post about refugees on the AfD Facebook page – a proxy of refugee salience on social media – in municipalities with one additional Facebook user.<sup>16</sup> We use three measures of social media usage:  $AfD Users/Pop.$ ,  $Nutella Users/Pop.$ , and  $I_{Many Nutella Users}$ . Recall from section Section 2 that these are highly correlated. Because refugee attacks are a function of the number of refugees in a given area, we scale the the attacks accordingly. In the robustness section, we explore a plethora of different variable transformations and show our results are highly robust.

This framework has three advantages. First, it makes reverse causality an unlikely concern, since in our setting the municipalities from which hate speech emanates are not the municipalities which are heavily exposed to social media. To illustrate this, the correlation between the total number of refugee posts in a municipality and the share of AfD Facebook users is only 0.07 (it is 0.02 for the Nutella user share and the Nutella dummy). This implies that municipalities are not disproportionately exposed to hate speech that originates in the same municipality. Furthermore, one out of the 4,466 municipalities cannot induce meaningful variation in the Germany-wide weekly  $Refugee Posts_w$  measure.<sup>17</sup> As a

<sup>16</sup>We also experimented with a host of alternative measures of refugee salience and right-wing social media usage (see section 3.6 and table A.20 in the online appendix).

<sup>17</sup>Note that using posts from *all* municipalities creates a slight difference between our strategy and a standard Bartik-type regression. In a standard Bartik regression the overall industry growth rate is weighted sum of the growth rates in the individual geographical areas and hence one uses a leave-one-out measure of industry growth rates. Due to the differences

result, our refugee salience measure is plausibly exogenous to each municipality. Second, the ratio of social media users is time-invariant and thus not caused by whether a municipality experiences a refugee attack in a particular week.<sup>18</sup> Third, the panel format allows us to abstract from unobserved weekly factors and municipality-level predictors of attacks using a full set of fixed effects.<sup>19</sup> In this panel framework, the week fixed effects absorb any changes in the number of anti-refugee incidents that affect all municipalities to the same extent, e.g. nationwide news events on increases in the number of refugees. The municipality fixed effects control for permanent differences in the number of anti-refugee incidents across municipalities, e.g. due to a stronger right-wing presence. In the online appendix, we show that the findings estimating Equation (1) also hold on the county- and state-level.

The required identifying assumption is that – conditional on covariates – hate crimes would have reacted similarly to shocks to refugee salience on Facebook in the absence of differences in local social media use. A concern for our empirical strategy is that social media usage is correlated with differences in other municipality characteristics that can explain the influence of the salience measure on refugee attacks. One might worry, for example, that municipalities with a higher share of AfD users might have a history of right-wing violence, which could be an omitted factor for the transmission between online hate speech and anti-refugee attacks.

We explore the validity of this concern for observable characteristics empirically. To do so, we regress our three measures of local Facebook penetration on a host of municipality characteristics (one-by-one) and plot the results in Table 2.<sup>20</sup> GDP per worker and the number of internet outages per capita are uncorrelated with all three proxies for social media access. Only a few other variables show a consistent pattern. For the AfD-based measure, we find a few significant correlations that are to be expected: consistent with a higher “demand” for online hate, there are more AfD users in municipalities with a higher ratio of non-Christians per capita (although the opposite holds for the share of immigrants). They also tend to concentrate in towns with older populations; a lower employment share of manufacturing industries; a higher share of people on welfare; and lower voter turnout. Unsurprisingly, we also find a positive correlation with the vote share for the AfD in the 2017 election. Municipalities with many AfD users on Facebook are also more likely to be the scene of

---

in our setting (described above), it is not crucial that we use a leave-one out salience measure. In an additional robustness check we nonetheless show that constructing the leave-one-out salience measure using a sub-set of geo-located posts leads to almost equivalent results.

<sup>18</sup>In the robustness section below, we alternatively measure local social media penetration before the start of the refugee crisis, at the cost of reducing the number of users we have location data for. It turns out that this adjustment makes little difference for the results.

<sup>19</sup>Note that the non-interacted terms for refugee posts and local users are absorbed by the fixed effects.

<sup>20</sup>Note that we standardize all variables by their standard deviation and mean to make the magnitudes comparable.

murders by neo-Nazis, although they appear to have less of a history of anti-semitic violence during the Third Reich.<sup>21</sup> Consistent with the evidence reported in Heintze (2017), there is a positive correlation with the share of high-school graduates, who are more likely to use social media in Germany.

Crucially, *all* of these correlations switch sign or become clearly statistically insignificant for our measures based on Nutella users, in particular the “Many Nutella Users” dummy. The remaining correlations that we find for the Nutella dummy measure are consistent with the general usage pattern of social media in Germany. More concretely, we find that people in the age groups from 18-24 and 25-34 years are more likely to use social media (Destatis, 2017), although only marginally: to illustrate, the average age in towns with many Nutella users is 44.9, compared to 45 in those with fewer users. According to Hölig and Hasebrink (2016), people in the age group between 18 and 34 years are also twice as likely to use social media as their main news source; as a result, they are less willing to pay for newspaper subscriptions (Hölig and Hasebrink, 2017). This likely explains the negative correlation between our social media measures with newspaper subscriptions. Finally, we find more Facebook users per capita in urban areas and larger cities, a result that has been widely documented in surveys on social media demographics (e.g. Pew Research Center, 2016). Importantly, we find no systematic correlation of social media with general internet usage. In summary, these findings suggest a clear pattern: the few municipality characteristics that vary with Facebook usage are those that one would expect if our measures reflect local differences in social media usage. Intuitively, it is less obvious why Nutella Facebook usage should be correlated with anti-refugee incidents, except through higher exposure to refugee salience on social media. Put differently, Nutella usage provides us with variation in exposure to anti-refugee sentiment on social media that is unlikely to be driven by higher observable “demand” for online hate speech or a higher observable tendency to commit hate anti-refugee crime. Table A.7 confirms this pattern using additional analysis based on median sample splits.

We additionally report results that include interactions of our anti-refugee sentiment measure with all of the variables listed in Table 2 to control for increases in anti-refugee incidents that can be explained by observable differences between municipalities. Controlling for these interactions does not substantially affect the size of the estimated coefficient of  $\beta$ .

To rule out the influence of potential unobservable factors, we exploit quasi-experimental variation in internet and Facebook outages as exogenous shifts in social media access. These disruptions reduce the access of people to social media while leaving the Germany wide refugee salience and local

---

<sup>21</sup>Note that we use scaled versions of the number of deported jews and letters to “Der Stürmer” from Voigtlander and Voth (2012) for the cross-sectional regressions shown here. The reason is that the availability of these variables is correlated with a city’s population ( $z = 7.73$  in a logit regression) and thus, mechanically, also with the social media measures.

characteristics unchanged. As such these outages enables us to identify the effect of social media under mild assumptions.

### 3.3 Main Results

Does social media have the potential to propagate hateful sentiment and cause violent incidents? Table 3 presents the main results from estimating Equation (1). The coefficient on the interaction of local Facebook usage and Germany-wide refugee posts is positive and highly significant in all specifications. In column (1), we start with the panel regressions without control variables. The coefficient of 0.533 implies a large economic effect of social media. Consider the case of 84 refugee posts per week, which is the mean of the variable (the median being 83). What does our estimate imply for the transmission of a typical degree of refugee salience depending on local AfD users? As a case study, consider the cities of Frankfurt/Main and Dresden, which are about one standard deviation apart at the AfD users over population measure. The estimated effect of a typical number of AfD refugee posts in a city like Dresden are 0.043 attacks per 10,000 asylum seekers, while it is 0.029 for a city such as Frankfurt.<sup>22</sup> This shift in the share of right-wing social media users implies around 50% more attacks on refugees. The implied effects are large, given that the average number of attacks on refugees in the sample is 0.038 per 10,000 asylum seekers.<sup>23</sup>

As discussed in Section 3.2 the population share of AfD Facebook users is correlated with observable municipality characteristics. Column (2) thus allows for the set of baseline control variables described above to enter, which we interact with the refugee posts. This decreases the point estimate by about 30%, but it does not alter the statistical significance of the result. For brevity, we report the full set of estimated coefficients in the online appendix (Table A.10). Strikingly, the interaction of refugee posts with the AfD vote share in the 2017 federal election is not significant. This indicates a difference between our measure of AfD Facebook users and general support for the AfD. As it turns out, only the coefficients for the interactions with GDP/worker, population density, share of population with a high school diploma (Abitur), and broadband access are statistically significant at conventional levels.

We next replace the share of AfD users with the share of people active on the Nutella Facebook page. Recall from Table 2 that this share is either uncorrelated or negatively correlated with measures

---

<sup>22</sup>To see this, one simply has to multiply the share of the population using the AfD page (e.g. 9.5 per 10,000 inhabitant for the case of Dresden) with the coefficient and the median number of posts, i.e.  $0.533 \times \frac{84}{10,000} \times 9.5 \approx 0.043$ . The calculation for Frankfurt is  $0.533 \times \frac{84}{10,000} \times 6.4 \approx 0.029$ .

<sup>23</sup>Note that the small values do not mean that the number of attacks is quantitatively irrelevant. Rather, it is an artifact of the large cross-section of 4,466 municipalities combined with the 111 week-long sample.

**Table 2: BALANCEDNESS OF SOCIAL MEDIA USAGE - CORRELATIONS**

	AfD Users/Pop.		Nutella Users/Pop.		<i>I</i> <sub>Many Nutella Users</sub>	
	$\hat{\beta}$	S.E.	$\hat{\beta}$	S.E.	$\hat{\beta}$	S.E.
<b>Baseline Controls</b>						
Population (2015) <sup>†</sup>	0.0362***	(0.0094)	0.0401***	(0.0069)	0.2236***	(0.0364)
GDP/Worker	-0.0236	(0.0496)	0.0404	(0.0462)	0.0481	(0.0305)
Population density	0.0362***	(0.0077)	0.0437***	(0.0072)	0.2360***	(0.0323)
AfD vote share (2017)	0.0430***	(0.0094)	-0.0293***	(0.0081)	-0.0467	(0.0301)
Share Abitur	0.0240**	(0.0121)	0.0439***	(0.0115)	-0.0347	(0.0303)
Share broadband access	-0.0098*	(0.0054)	0.0098	(0.0073)	0.0497	(0.0307)
Share immigrants	-0.0182***	(0.0049)	0.0258***	(0.0078)	0.0482	(0.0315)
<b>Additional Media and Internet Controls<sup>†</sup></b>						
Internet outages/Pop.	-0.0082	(0.0051)	-0.0004	(0.0055)	0.0333	(0.0339)
Registered domains/Pop.	0.1258***	(0.0431)	0.0697*	(0.0403)	0.0472	(0.0312)
News paper sales/Pop.	-0.0371***	(0.0092)	-0.0544***	(0.0109)	-0.1269***	(0.0354)
<b>Additional Right-Wing Controls</b>						
Nazi murders/Pop. <sup>†</sup>	0.0201***	(0.0069)	0.0101	(0.0091)	0.0292	(0.0294)
NPD vote share (2017)	0.0439***	(0.0095)	-0.0046	(0.0076)	-0.0643**	(0.0301)
Deported Jews/Jews (1933)	-0.0269**	(0.0110)	0.0019	(0.0157)	0.0007	(0.0671)
Stürmer letters/Pop. (1933)	-0.0217*	(0.0127)	-0.0395**	(0.0161)	-0.0345	(0.0619)
<b>Additional Socio-Economic Controls</b>						
Average age	0.0547***	(0.0125)	0.0248***	(0.0085)	-0.0541*	(0.0308)
Share benefit recipients	0.0236***	(0.0080)	0.0306***	(0.0079)	-0.0181	(0.0305)
Share non-Christians	0.0606***	(0.0133)	0.0336***	(0.0087)	-0.0266	(0.0308)
Manufacturing share (%)	-0.3548***	(0.0454)	-0.1440**	(0.0567)	-0.0309	(0.0303)
<b>Additional Voting Controls (2017 Election)</b>						
CDU vote share	-0.0513***	(0.0111)	-0.0246***	(0.0079)	0.0245	(0.0305)
SPD vote share	-0.0047	(0.0048)	0.0696***	(0.0096)	0.0026	(0.0301)
Left vote share	0.0510***	(0.0108)	0.0038	(0.0072)	-0.0271	(0.0304)
Green vote share	-0.0366***	(0.0089)	-0.0098	(0.0071)	0.0352	(0.0304)
FDP vote share	-0.0175***	(0.0053)	0.0077	(0.0068)	-0.0072	(0.0302)
Pirate vote share	-0.0115**	(0.0054)	0.0252***	(0.0065)	0.0414	(0.0304)
Voter turnout	-0.0452***	(0.0107)	-0.0283***	(0.0075)	0.0028	(0.0304)
<b>Additional Demographic Controls</b>						
Share aged 3-14	-0.0572***	(0.0117)	-0.0177**	(0.0073)	-0.0169	(0.0300)
Share aged 15-17	-0.0661***	(0.0127)	-0.0100	(0.0074)	-0.0230	(0.0302)
Share aged 18-24	-0.0194***	(0.0068)	0.0158**	(0.0067)	0.1441***	(0.0301)
Share aged 25-29	0.0179***	(0.0057)	0.0319***	(0.0066)	0.2466***	(0.0307)
Share aged 30-39	-0.0185***	(0.0067)	-0.0046	(0.0067)	0.0877***	(0.0300)
Share aged 40-49	-0.0365***	(0.0083)	-0.0170**	(0.0066)	-0.0417	(0.0298)
Share aged 50-64	-0.0011	(0.0051)	0.0020	(0.0067)	-0.0191	(0.0299)
Share above 65	0.0347***	(0.0094)	0.0283***	(0.0077)	0.0428	(0.0301)

*Notes:* This table reports estimated coefficients from regressing the variables in the left column (one-by-one) on the social media measures listed above. All coefficients are standardized to have a mean of 0 and standard deviation of 1. Robust standard errors are in parentheses. The variables GDP/Worker, Registered domains/Pop., and Manufacturing share are estimated on the county-level; all other variables on the municipality-level. The "Many Nutella Users" dummy is 1 for municipalities in the top tercile of Nutella users per capita within a county, and 0 otherwise. See text for discussion.

**Table 3: MAIN RESULTS: FACEBOOK POSTS AND HATE CRIME**

Dependent Variable	Refugee Attacks/Refugees					
	AfD Users/Pop.		Nutella Users/Pop.		$I_{Many\ Nutella\ Users}$	
Social Media Measure	(1)	(2)	(3)	(4)	(5)	(6)
Social media users/Pop. $\times$ Refugee posts	0.533*** (0.151)	0.369*** (0.113)	0.428*** (0.119)	0.287*** (0.091)	2.884*** (0.675)	2.317*** (0.621)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts		Yes		Yes		Yes
Observations	495,726	480,963	495,726	480,963	495,615	480,852
Number of municipalities	4,466	4,333	4,466	4,333	4,465	4,332
$R^2$	0.045	0.045	0.045	0.045	0.045	0.045

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on local social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* and *Nutella users/Pop.* are the ratio of people with any activity on the Facebook pages of the AfD and Nutella, respectively, scaled over the total municipality population. The *Many Nutella Users Dummy* is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise. *Refugee posts* is the Germany-wide number of posts on the AfD's Facebook wall containing the word refugee ("Flüchtling"), divided by 10,000 for readability. See text for an explanation of the control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, \*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

of local right-wing presence. The results are presented in columns (3) and (4). The coefficients on the interaction term here are very similar, and still highly statistically significant. The inclusion of the interacted baseline controls in column (4) again reduces the effect by about 30%. In order to abstract from the remaining observable municipality differences we documented before, columns (5) and (6) use the top tercile of *Nutella users/Pop.* to proxy for social media usage. We find comparable estimates; perhaps unsurprisingly, the inclusion of controls has a much smaller effect here. The dummy also has the advantage that we can easily “read off” which effect local propagation of right-wing social media has: the coefficient of around 2.8 implies that, even within the same county, a municipality with many Facebook users has approximately 0.024 more refugee attacks (per 10,000 refugees) than a municipality with few users in a typical week.<sup>24</sup> This corresponds to an increase of almost two-thirds of the mean of the dependent variable, a large effect.

We further address the concern of observable differences between municipalities with a higher share of AfD Facebook users by introducing a richer set of controls in Table 4.<sup>25</sup> For convenience, column (1) reproduces the estimation with the baseline controls to aid a comparison of the coefficients. In each of the following columns, we add a new control set plotted in Table 1 and Table 2, which we interact with our measure of anti-refugee sentiment.

In column (2), we first include additional controls right-wing controls, namely the vote share of the neo-Nazi party NPD and proxies for anti-minority violence. Recall that the correlation of AfD usage with these variables showed a somewhat mixed pattern. Strikingly, the interacted controls have hardly any effect on the estimated coefficient. This suggests that our findings cannot be easily explained by a higher presence of neo-Nazis or far-right groups in municipalities with more AfD Facebook users. Next, we investigate if our effect is driven by general media exposure in column (3). We again find almost no change in the point estimate and thus no evidence that the increase in attacks is driven by stronger exposure to the news cycle either through traditional media consumption or general internet affinity. Adding more socio-economic controls in column (4) or controls for the vote share of all other major parties in the 2017 election in column (5) has similarly little effect.<sup>26</sup>

The only noticeable effect on our results we uncover comes from the inclusion of more flexible controls for the age structure of each municipality (column (6)). The reduction in the coefficient can likely be explained by the fact that social media use is concentrated in particular age groups,

<sup>24</sup>The calculation is based on the mean number of refugee posts (84) as above, i.e.  $2.8 \times \frac{84}{10,000} \approx 0.024$ .

<sup>25</sup>In the online appendix Table A.12, we reproduce these regressions using our Nutella measures of social media usage.

<sup>26</sup>In unreported regressions, we further find that the results are nearly identical if we control for the election results of the 2013 federal election instead.

which we discussed in Section 3.2. Column (7) shows that the results are also robust to including all the interacted controls at the same time. Finally, in column (8), we report an even more stringent specification including *county*  $\times$  *week* fixed effects in the regression. Essentially, this means comparing neighboring municipality outcomes in the *same county* in the *same week*.<sup>27</sup> Unsurprisingly, this decreases our point estimate by a substantial margin, but the effect remains statistically significant at the 5% level. This suggests that the observed effect of municipality-level social media exposure is not driven by unobservable county differences.

We explore effect heterogeneity across different type of anti-refugee incidents in Appendix Table A.11. We find that the effects are entirely driven by more violent crimes, namely arsons, assaults, and miscellaneous property damage. We do not find significant transmission effects on demonstrations (or the few suspected cases in the sample). This is consistent with the notion that online hate speech can act as a propagating mechanism for violent crimes in particular. It also contrasts with the findings in Enikolopov et al. (2016), who find that social media spurs protest participation in Russia.

Overall, the findings we present in this section suggest that exposure to right-wing refugee salience on social media is a predictor of violent attacks on refugees. This is true both for municipalities with many right-wing Facebook users as well as those with high social media affinity that is unrelated to observable municipality characteristics. In the next sections, we attempt to narrow down channels that are specific to hate transmission via right-wing social media and subject our baseline findings to a number of robustness exercises.

### 3.4 Quasi-Experimental Evidence from Internet Outages

To further isolate the importance of social media and rule out that our results are driven by unobservables, we next draw on internet and Facebook outages as sources of quasi-experimental variation. During our sample period, around 50% of worldwide Facebook users accessed the platform with their computer; in Germany, this share is likely to be higher because of the relatively slow adaption of mobile internet.<sup>28</sup> This exposes a relatively large share of social media users to disruptions in their access to the internet infrastructure.

We start by exploiting user-reported internet disruptions to assess whether the connection between social media and hate crime we have established are indeed due to online activity. As

---

<sup>27</sup>Note that this leads to a slight drop in observations, because some counties only contain a single municipality (e.g. major cities).

<sup>28</sup>Data on Facebook usage patterns reported on Statista.com and on mobile internet usage in Germany on (also on Statista.com) support this assessment.

**Table 4: MAIN RESULTS WITH ADDITIONAL CONTROL VARIABLES**

	Additional interacted controls							
	Baseline controls only (1)	Right-wing controls (2)	Media controls (3)	Socio-economic controls (4)	2017 vote controls (5)	Age structure controls (6)	All controls (7)	County $\times$ Week FE (8)
AfD users/Pop. $\times$ Refugee posts	0.369*** (0.113)	0.360*** (0.111)	0.363** (0.144)	0.355*** (0.112)	0.346*** (0.108)	0.223*** (0.083)	0.208** (0.100)	0.151*** (0.060)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	–
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Right-wing controls (4) $\times$ Posts		Yes					Yes	Yes
Media controls (3) $\times$ Posts			Yes				Yes	Yes
Socio-econ. controls (4) $\times$ Posts				Yes			Yes	Yes
Election controls (7) $\times$ Posts					Yes		Yes	Yes
Age controls (8) $\times$ Posts						Yes	Yes	Yes
County $\times$ Week FE								Yes
Observations	480,963	480,963	358,308	475,302	480,963	480,963	354,423	483,849
Number of municipalities	4,333	4,333	3,228	4,282	4,333	4,333	3,193	4,359
$R^2$	0.045	0.045	0.043	0.046	0.045	0.046	0.043	0.111

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on local social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* are the ratio of people with any activity on the Facebook page of the AfD, scaled over the total municipality population. *Refugee posts* is the Germany-wide number of posts on the AfD's Facebook wall containing the word refugee ("Flüchtling"), divided by 10,000 for readability. See text and table Table 2 for the included control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

described above, the data on internet outages we use do not only measure disruptions resulting from physical infrastructure damages, but also provider-specific problems; in fact, many of the large outages in our data that received media attention originated from problems at a single provider. These outages are also widely geographically dispersed: Figure A.4 visualizes the distribution of the number of disruptions per capita across Germany. We define our baseline measure as an internet outage that is above the 75th percentile of this variable (see section Section 2 for more details). This gives us 313 municipality-week observations that are associated with severe internet problems. In the online appendix, we show the robustness of our results to alternative definitions or using the continuous ratio of the local disruptions per capita with over 1000 internet problems (see Table A.14). Crucially, the frequency of internet problems is virtually uncorrelated with all of our social media measures (see Table 2); as such, they provide exogenous variation that is not already captured by our variables on local Facebook usage. The number of internet problems is also orthogonal to the total number of refugee attacks in a given municipality ( $t - stat = -1.33$ ).

Does our measure of local internet outages indeed have an impact on the activity of affected users on the AfD's Facebook page? We address this question using a sub-sample of geo-coded measures of social media activity, namely the number of refugee posts as well as total posts, likes, comments, and shares on the AfD page. In particular, table Table 5 reports the mean number of these variables, divided by whether a municipality experiences an outage in a given week.<sup>29</sup> We also report the associated  $t$ -statistic under the null hypothesis of no difference in means. In all cases, we observe significantly lower local activity on Facebook for municipalities that experience an internet disruption. This suggests that the decrease in social media activity due to such disruptions is not made up by users accessing Facebook with their mobile phones; the outages thus also decrease exposure to right-wing social media content in the affected areas.

If there is indeed a role for social media in inciting anti-refugee incidents, we would expect that reduced exposure to social media due to internet outages in a municipality would have a mediating effect on the number of hate crimes. We test this hypothesis by interacting the main effect of interest  $Refugee Posts_w \times Socialmedia Users/Pop_m$  with  $Internet Problems_{mw}$ , our dummy for large internet outages. The regression results in columns (1) to (3) of Table 6 suggest that internet problems reduce the impact of social media on anti-refugee violence. While the main effect of refugee posts and

---

<sup>29</sup>Because we only have a smaller number of Facebook actions that we can directly tie to local users, we use a more conservative cut-off for internet outages here: all those above the *median* in the ratio of user-reported internet disruptions to population (compared to the top quartile cut-off used in the baseline regression). In the online appendix, we show that these findings also hold for our baseline measure and conduct further robustness checks (see Table A.13). They also hold *within* a given municipality, but this again reduces the effective sample size (unreported).

**Table 5: DO LOCAL INTERNET OUTAGES REDUCE LOCAL FACEBOOK ACTIVITY?**

	Refugee posts	Total posts	Total likes	Total comments	Total shares
No outage	0.008	0.069	0.153	0.071	0.157
Internet outage	0.003	0.003	0.000	0.000	0.000
Diff. > 0?	1.466*	16.397***	26.264***	29.272***	4.654***

*Notes:* This table presents the arithmetic mean of local measures of Facebook activity based on linking user locations with posts, likes, comments, and shares. “Internet outage” are municipality-week observations in which a local internet outage occurs. The row “Diff. > 0” presents the results of one-sided t-tests under the null hypothesis that the mean values are larger in the “No outage” sample, and the associated p-values (we do not assume equal variances across samples).

social media exposure is similar to our baseline coefficient, the triple interaction term is negative and statistically significant for all three measures of local Facebook usage.

Quantitatively, internet outages appear to mitigate the entire effect of social media. The size of the triple interaction coefficient is larger than the baseline coefficient. This indicates that in weeks of high refugee salience, the likelihood of an anti-refugee incident is smaller in municipalities with high Facebook usage that experience an internet outage than in municipalities with low Facebook usage *without* an outage. This finding seems intuitive, since we focus on major disruptions that affect a significant fraction of the local population. Internet outages effectively cut off some users from right-wing social media content. These outages should thus induce a larger effect on hate crimes than municipality differences in local social media usage (captured by the baseline coefficient).

It is also worth pointing out that the number of internet outages itself does not have any significant effect on the number of anti-refugee incidents. This suggests that our findings are not driven by potential displacement effects of time usage, e.g. time spent fixing the internet, since these effects should be independent of the extent of Facebook usage or anti-refugee sentiment.

Could it be that we observe reduced hate crimes because users are cut off from the internet generally, and not from social media in particular? We believe this is unlikely to be the case for two reasons. First, our social media use proxies are essentially orthogonal to measures of general internet affinity. The triple interaction term thus captures the differential effect of users being cut off from social media and not the internet in general. Second, when we control for interactions of internet disruptions with our measures of internet use (broadband access and per capita internet domains), the coefficients of the triple interactions are statistically insignificant; in fact, the inclusion of these

additional variables *increases* the significance of the main effects.<sup>30</sup>

To rule out that our results are spurious, we additionally perform a randomization test. Instead of the actual internet disruptions, we randomly define 313 placebo outages in the municipalities that experience at least a single internet outage during our sample period. We then estimate the same regression using 500 sets of placebo outages to get a sense of how likely it is to accidentally find significant values. Only in 3.2% of the cases we find coefficients that are significant at the 5% level. This suggests that it is unlikely that our findings are purely coincidental. The results are similar if we randomize the internet outages over all municipalities. In online appendix Figure A.6 we show the full distribution of  $t$ -stats from both randomization tests.

To dig deeper, we use eight major Germany-wide Facebook outages as an additional source of exogenous variation in social media access. Table A.4 outlines the details of each of these and links to relevant press reports. The advantage of these is that they are Facebook-specific and therefore do not affect other potential channels of online hate speech transmission. In fact, we find that they are uncorrelated with the total number of weekly internet disruptions in the time series of 111 weeks ( $t = -0.28$ , unreported). The disadvantage is that they affect users nationwide, and sometimes globally, and thus only vary by time and not by municipality. We present the results of interacting these variables analogous to the internet outages in columns (4) to (6) of Table 6.

The results reveal a striking pattern. While the Facebook outages only vary by week, the triple interaction term with social media usage and our refugee salience measure attracts a statistically significant negative estimate in all three specifications. Quantitatively, we find that Facebook disruptions *fully* undo the effect of social media propagation. For example, consider that the coefficient of *Nutella Users/Pop.* and *Refugee Posts* is 0.317 in column (5), while that on the triple interaction is  $-0.337$ .

Taken together, the evidence presented in this section points towards the interpretation that our results are indeed driven by the propagation of hateful sentiment through social media.

### 3.5 Additional Results

In the following subsections, we present an additional set of test for our hypothesis that online hate speech propagated through social media can amplify refugee attacks.

---

<sup>30</sup>These results are available upon request.

**Table 6: QUASI-EXPERIMENTAL EVIDENCE: INTERNET AND FACEBOOK OUTAGES**

	Local Internet Outages			Country-Wide Facebook Outages		
	AfD Users/Pop. (1)	Nutella Users/Pop. (2)	<i>IMany Nutella Users</i> (3)	AfD Users/Pop. (4)	Nutella Users/Pop. (5)	<i>IMany Nutella Users</i> (6)
Outage	-0.029 (0.019)	-0.019 (0.026)	-0.027 (0.025)	-	-	-
Outage × Posts	2.467 (1.922)	-0.092 (2.144)	-0.019 (2.065)	-	-	-
Social media users/Pop. × Outage	0.007** (0.003)	0.005 (0.007)	0.039 (0.027)	0.000 (0.001)	0.000 (0.000)	0.001 (0.009)
Social media users/Pop. × Refugee posts	0.369*** (0.113)	0.287*** (0.091)	2.321*** (0.622)	0.398*** (0.121)	0.317*** (0.098)	2.588*** (0.664)
Social media users/Pop. × Posts × Outage	-2.037** (0.825)	-1.108* (0.662)	-5.160** (2.587)	-0.327* (0.185)	-0.337** (0.172)	-2.991*** (1.049)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) × Posts	Yes	Yes	Yes	Yes	Yes	Yes
Observations	480,963	480,963	480,852	480,852	480,963	480,963
Number of municipalities	4,333	4,333	4,332	4,333	4,333	4,332
$R^2$	0.045	0.045	0.045	0.045	0.045	0.045

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media usage. The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* and *Nutella users/Pop.* are the ratio of people with any activity on the Facebook pages of the AfD and Nutella, respectively, scaled over the total municipality population. The *Many Nutella Users Dummy* is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise. *Refugee posts* is the Germany-wide number of posts on the AfD's Facebook wall containing the word refugee ("Flüchtling"), divided by 10,000 for readability. Internet outages are defined as municipality-weeks that are in the top quartile of the ratio of reported internet outages to population. Facebook outages refer to weeks where Facebook experienced considerable disruptions; note that *Outage* and *Outage × Posts* are absorbed by the week fixed effects. See text and the online appendix for details. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

### 3.5.1 Placebo Tests for other Posts on the AfD Facebook page

If the channel we are uncovering is indeed specific to refugees, we should expect that refugee attacks should show much weaker correlations with posts on other topics on the AfD Facebook page. We test this hypothesis formally in Table 7, where we plot the baseline estimation with refugee posts in column (1) for convenience. We also report coefficients for standardized post measures (with mean 0 and standard deviation of 1) in square brackets to allow for a comparison of effect sizes across the different posts. Next, we estimate Equation (1) using all posts except those containing the word *refugee* (“Flüchtling”) in column (2). The estimate is statistically indistinguishable from zero. We also repeat our baseline test using posts containing the words “Muslim”, “Islam”, “Jude” (Jew), or “EU” – the latter is motivated by the AfD’s long-standing criticism of the European Union. For none of the other measures, we find a significant relationship between the number of post and the number of attacks; all estimated coefficients are considerably smaller in standardized terms compared to the baseline measure. This shows the specificity of our anti-refugee measure: the effect we are capturing does not appear to be an artifact of general anti-minority sentiment, but rather a predictable result of increased animosities towards refugees on social media in particular weeks.

**Table 7: NON-REFUGEE FACEBOOK POSTS AND HATE CRIMES**

	Refugee (Baseline) (1)	All other posts (2)	Muslim posts (3)	Islam posts (4)	Jew posts (5)	EU posts (6)
AfD users/Pop. × FB Posts	0.369*** [2.264] (0.113)	0.016 [0.766] (0.011)	0.259 [0.470] (0.252)	-0.058 [-0.199] (0.118)	2.758 [0.808] (2.234)	-0.023 [-0.046] (0.164)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) × Posts	Yes	Yes	Yes	Yes	Yes	Yes
Observations	480,963	480,963	480,963	480,963	480,963	480,963
Number of municipalities	4,333	4,333	4,333	4,333	4,333	4,333
$R^2$	0.045	0.045	0.045	0.045	0.045	0.045

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on local social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on the AfD’s Facebook page to population. *FB Posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the words in the top row, divided by 10,000 for readability, with the baseline being *refugee* (“Flüchtling”). Standardized coefficients reported in square brackets, based on variable transformations with a mean of 0 and a standard deviation of 1. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

### 3.5.2 Results for the Intensive Margin of Facebook Usage

If social media works as a propagating mechanism for hate speech, we would also expect that the effect size increases with the local usage intensity of AfD's Facebook page, since a higher intensity goes along with a higher exposure to hate speech. In the case of Facebook, one way to measure this intensity is to look at how frequently the average AfD user in a municipality posts on the AfD wall or how many likes or comments she receives on her own posts. If we are indeed capturing social network effects, we would expect that anti-refugee sentiment has particularly strong effects in areas where the average AfD user shows a higher level of social media engagement.

We explore this issue empirically in Table 8, where we interact our main interaction term with of the total number of local posts on the AfD wall, as well as the number of comments and likes on AfD posts, all scaled over the number of AfD users in a municipality.<sup>31</sup> Note that these user intensity measures are not systematically correlated with local Facebook penetration, city size, and population density. As such, they create additional variation on the municipality-level that is mostly orthogonal to the correlates of social media usage we outlined in Section 3.2.

The results suggest that the intensity of local Facebook usage matters: all three triple interaction terms are positive and statistically significant. Consistent with the hypothesis that social media enables hateful sentiment to spread, higher user activity increases the effect of the Germany-wide transmission of refugee salience. Importantly, these effects work on top of our baseline interaction term, which remains similar in magnitude and highly statistically significant throughout. To put things into perspective, consider the smallest effect we find here, which is for the average number of likes per AfD user (column (3)). The coefficient on the triple interaction term of 0.013 implies that a one standard deviation increase in the likes per user (around 12) increases the baseline effect by around one third.<sup>32</sup> Note that we find significantly negative estimates for the double interaction of the AfD user variable and our reach measures. While we are not interested in these per se, their negative effect is outweighed by the positive coefficient of the triple interaction in any week with more than three refugee posts (the median being 83).

---

<sup>31</sup>Note that we can only construct these measures on the intensive margin of municipalities where we can identify at least a single AfD user. Our main results also hold in this sub-sample, which we show in Table A.19 in the online appendix.

<sup>32</sup>To see this, consider that the total effect including interaction is calculated as  $0.317 + 0.013 \times 12 \approx 0.47$ , which is about 1/3 larger than the baseline effect of 0.369.

**Table 8: SOCIAL MEDIA REACH AND HATE CRIME PROPAGATION**

Reach (per AfD User)	Number of posts (1)	Received comments (2)	Received Likes (3)
AfD users/Pop. $\times$ Refugee posts	0.308** (0.132)	0.304** (0.133)	0.317** (0.132)
AfD users/Pop. $\times$ Posts $\times$ Reach	0.056** (0.022)	0.032** (0.013)	0.013** (0.006)
Refugee posts $\times$ Reach	-0.179** (0.081)	-0.068** (0.029)	-0.033* (0.018)
Week FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts	Yes	Yes	Yes
Observations	381,840	381,840	381,840
Number of Municipalities	3,440	3,440	3,440
$R^2$	0.046	0.046	0.046
Corr(Measure,Population)	0.012	0.012	0.010
Corr(Measure,Population density)	0.025	0.038**	0.026
Corr(Measure,AfD users/Pop.)	0.016	0.005	0.019
Corr(Measure,Nutella users/Pop.)	-0.010	0.001	-0.005
Corr(Measure,Average age)	-0.019*	-0.011	-0.021*

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on local social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. The reach variables in the top row refer to the number of local posts on the AfD wall, as well as comments and likes for AfD posts, all scaled by the number of AfD users (municipalities with zero users are dropped). See text for an explanation of the control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

### 3.5.3 Distracting News Events and Social Media

As an additional piece of analysis, we investigate the effects of news shocks on the transmission of online hate speech to real-world actions. Durante and Zhuravskaya (2016) provide evidence for the hypothesis that the Israeli army is more likely to strike against Palestinian targets when US media outlets are distracted by other events. In our case, we hypothesize that other important news events might distract people from the so-called refugee crisis. This is somewhat analogous to our Facebook outages in that we exploit additional exogenous weekly variation, albeit with a different goal: if we are correct that our refugee salience measure captures public discourse, we should see a role for news distractions.

To measure these news shocks, we obtain data from Google Trends on the weekly search interest on the terms “Brexit”, “Trump”, and “UEFA Euro 2016”. Including these Google trends as a further interaction in our regressions allows us to investigate whether these events crowd-out the salience of refugees (see Table 9). More precisely, we would expect that these events *decrease* the effect of social media transmission on refugee attacks. We focus on the *AfD Users/Pop.* measure for brevity, but the results are similar for the Nutella-based social media proxies (unreported).

For each of the events in columns (1), (2), and (3), we find a significant negative coefficient on the number of anti-refugee incidents for the triple interaction with distracting news. The negative sign of the coefficient indicates that during weeks of major news events increases in refugee salience have a smaller effect on anti-refugee incidents. This finding is consistent with the salience of other events crowding out that of refugees, which reduces hate crimes in municipalities with more AfD social media users.

## 3.6 Exploring Alternative Measures of Refugee Salience

Our baseline findings are based on a measure of refugee salience on right-wing social media constructed from the AfD’s Facebook page. In this section, we further compare our measure with salience measures constructed from newspaper articles and right-wing protest participation. This exercise serves three purposes. First, we want to understand to what extent our Facebook measure accurately reflects general refugee salience in Germany. Second, we want to understand whether social media transmission has different effects depending on the type of sentiment it propagates. Third, we are interested in whether our social media-based measure provides information about the likelihood of anti-refugee incidents *over and above* that contained in the news media or right-wing demonstrations. To this end, recall that it is the variation in *local Facebook exposure* to Germany-wide salience that allows us to infer a role

**Table 9: NEWS SHOCK SALIENCE AND HATE CRIME PROPAGATION**

	Distracting News		
	Brexit (1)	Trump (2)	Football (3)
AfD users/Pop. $\times$ Refugee posts	0.417*** (0.123)	0.594*** (0.167)	0.382*** (0.116)
AfD users/Pop. $\times$ Posts $\times$ News shock	-0.177** (0.082)	-0.068** (0.029)	-0.011* (0.006)
AfD users/Pop. $\times$ News shock	0.001** (0.001)	0.000* (0.000)	0.000 (0.000)
Week FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts	Yes	Yes	Yes
Observations	480,963	480,963	480,963
Number of municipalities	4,333	4,333	4,333
$R^2$	0.045	0.045	0.045

This table presents the estimated coefficients from a regression of hate crimes against refugees on local social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. The news shocks refer to the Google searches as indicated in the text. See text for an explanation of the control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01%, 0.05%, and 0.1%, respectively.

for social media, not the mere time series variation. Intuitively, we would expect to find similar results independent of source from which we construct the salience measure.

We investigate the effect of different salience measures on anti-refugee incidents by type in Table 10, where each cell represents a separate regression.<sup>33</sup> To aid the comparison with our main results, panel A plots the coefficients using our AfD refugee post measure. We also present standardized coefficients in square brackets (based on transformations of the salience measures with a mean of 0 and standard deviation of 1), which we use to compare effect magnitudes across the empirical specifications.

As a first step, we re-run our main regression Equation (1) using a salience measure based on refugee news coverage, interacted with the ratio of AfD users to population, in panel B. Our measure of news coverage is the first principal component of the number of articles on refugees in the major German newspapers (described in more detail in Section 2). As it turns out, this proxy has a high correlation of around 0.8 with the number of refugee posts on the AfD page (see Table A.15). This lends further credence to the idea that the number of anti-refugee post on the AfD's Facebook page is a sound proxy for the Germany-wide anti-refugee salience.

The interacted news coverage measure has a positive coefficient and is highly statistically significant in four out of the five dependent variables. In terms of magnitude, it fares fairly similar to the AfD-based sentiment for the overall number of incidents in column (1) and property damage in column (3), but has an around 30% lower effect on arsons in column (2). We find no evidence for an effect on assaults (column 4), but a much larger coefficient for local demonstrations in column (5). Compared to our Facebook post proxy, it thus appears that refugee salience in the media has a somewhat weaker predictive ability for violent incidents, but a stronger effect on non-violent protests.

Next, we consider news coverage in the “traditional” right-wing media outside of Facebook by using the interaction term of articles about refugees on the website of Compact Magazine, which is considered a voice for the AfD, in panel C. The number of Compact articles about refugees has a correlation of around 0.7 with the AfD posts. This proxy is also statistically significant with an around 20% lower magnitude than the social media measure. Right-wing media has no effect on protests (column 5), but has quantitatively meaningful correlations with property damages and assaults, similar to those of our Facebook proxy.

We further attempt to capture right-wing anti-refugee salience outside of media reports by

---

<sup>33</sup>For brevity, we limit the estimations here to the *AfD users/Pop.* measure. The results are highly similar using the Nutella-based proxies (available upon request).

**Table 10: EXPLORING ALTERNATIVE SENTIMENT MEASURES**

	Refugee Incident Type				
	Any (1)	Arson (2)	Property Damage (3)	Assault (4)	Demonstration (5)
<b>Panel A: Baseline results</b>					
AfD users/Pop. $\times$ Refugee posts	0.369*** [2.264] (0.113)	0.027** [0.164] (0.012)	0.227*** [1.391] (0.072)	0.078** [0.477] (0.032)	0.039 [0.238] (0.029)
<b>Panel B: Media coverage</b>					
AfD users/Pop. $\times$ Media coverage	1.069*** [2.050] (0.316)	0.063** [0.120] (0.030)	0.700*** [1.341] (0.220)	-0.002 [-0.000] (0.078)	0.317*** [0.607] (0.106)
<b>Panel C: Right-wing media coverage</b>					
AfD users/Pop. $\times$ Compact articles	0.629*** [1.863] (0.207)	0.028 [0.084] (0.020)	0.418*** [1.237] (0.138)	0.131* [0.388] (0.074)	0.052 [0.154] (0.043)
<b>Panel D: Right-wing protest participation</b>					
AfD users/Pop. $\times$ PEGIDA	0.390** [1.532] (0.186)	0.01 [0.040] (0.016)	0.092 [0.361] (0.058)	0.126 [0.496] (0.155)	0.163*** [0.642] (0.059)
Week FE	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts	Yes	Yes	Yes	Yes	Yes

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media usage as in Equation (1). Each cell represents a separate regression. The dependent variable is the ratio of municipality-level refugee incidents to the number of asylum seekers (listed in the top row); *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. *PEGIDA* is the weekly number of “PEGIDA” demonstrators in Dresden. *Media coverage* is the first principal component of the number of news items on major German newspaper websites containing the word refugee (“Flüchtling”). *Compact articles* is the equivalent measure based on the right-wing Compact Magazin’s website. See text for an explanation of the control variables. Standardized coefficients reported in square brackets, based on variable transformations with a mean of 0 and a standard deviation of 1. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

including the interaction of local Facebook usage with the number of PEGIDA demonstrators in column (5). The “Patriotic Europeans Against the Islamisation of the West” (*Patriotische Europäer gegen die Islamisierung des Abendlandes*) or PEGIDA movement in Dresden is a far-right anti-Islam demonstration with an explicit focus on refugees, which was founded in October 2014 and has held weekly demonstration ever since. The initiator of PEGIDA, as well as multiple speakers at their events, have been convicted of “incitement of the people” (*Volkshverhetzung*) – among other things for statements at their rallies such as “unfortunately, the concentration camps are currently not in use.” As such, we take the the number of PEGIDA participants as a reasonable proxy for right-wing anti-refugee salience that is not limited to social media. The PEGIDA measure might for example capture a delayed effect, since the participation in protests required more planning than Facebook posts. In the time series, the number of PEGIDA participants is virtually uncorrelated with refugee posts on the AfD page and news reports (see table A.15).

We find in panel D that *general* right-wing sentiment also predicts more refugee incidents in towns with higher right-wing social media usage. Intuitively, this result is fully driven by the effect of PEGIDA on municipality-level *demonstrations*, with clearly insignificant and small (standardized) coefficients for violent incidents. This contrasts with our AfD-based and newspaper measures, which predict violent crimes. Quantitatively, the coefficient in column (5) implies that a one-standard deviation shift in the PEGIDA count has an around three times larger effect on demonstrations compared to the AfD refugee posts. This suggests that social media can transmit different types of sentiments, which may spur both violent crimes or non-violent protests.

How do these different sentiment measures compare quantitatively when they are stacked against each other? To get a more direct comparison, Table A.16 in the online appendix runs a “horse race” of our social media salience measure against the others. The dependent variable is the number of refugee attacks to asylum seekers. Column (1) reproduces the paper’s baseline result. In column (2), we introduce the media coverage measure, which is now clearly insignificant. Our main interaction term decreases in size by approximately 20% but remains statistically significant just below the 5% threshold.<sup>34</sup> The drop in statistical power is unsurprising, given that the two sentiment variables are highly correlated. Quantitatively, the standardized coefficients now imply that a one-standard deviation shift in Facebook salience has a more than *twice* the effect of additional news coverage. This result

---

<sup>34</sup>Without the baseline controls included here, our main effect is still highly significant with a p-value of 0.016; the news coverage interaction is still clearly insignificant (available upon request). In the online appendix Table A.17, we replicate the results plotted here using the coverage on the individual newspapers’ websites, which yields highly similar results.

suggests that right-wing anti-refugee sentiment has an independent effect over and above the news cycle.

In column (3), we add the Compact articles interaction term, which is also insignificant, but leaves the coefficient size and statistical significance of the AfD posts nearly unchanged. The refugee post effect is almost four times that of the Compact coverage. At last, we introduce the PEGIDA right-wing sentiment proxy in column (4), which in fact *increases* the estimate of the refugee post interaction. In quantitative terms, the coefficients imply that the AfD Facebook proxy has a 40% larger effect than the PEGIDA measure when entered jointly. Including all other sentiment measures jointly in column (5) yields similar results.

Taken together, these findings indicate that the number of (anti-)refugee posts on the AfD's Facebook page is a solid measure of refugee salience. This social media-based measure also has predictive power for refugee attacks over and above news coverage. Facebook sentiment also appears to drive particularly violent incidents, while right-wing protest participation for PEGIDA is associated with more local demonstrations, i.e. non-violent incidents.

### **3.7 Additional Robustness Exercises**

The results in the previous sections suggest a tight link between right-wing online activity in social networks and violent hate crimes. We now subject this finding to a number of further robustness exercises, which we present in the online appendix for brevity. Throughout, our results remain remarkably stable.

First, we consider a dynamic specification with a lagged dependent variable in column (1) of Table A.18. In Figure 3a, we showed that there is some persistence in the time series of refugee incidents, which might introduce misspecification concerns. The AR(1) regression, however, yields a coefficient of 0.368 for our main effect, almost equivalent to the baseline result of 0.369. In column (2), we weight the regression by municipality population to check for the possibility that minor villages or population density are driving the result. However, we find that the result is hardly affected by this perturbation. The same holds true for column (3), where we replace the total number of AfD users by the pre-sample number of users to address potential endogeneity concerns of the page's popularity as a result of increased refugee violence.<sup>35</sup>

---

<sup>35</sup>As explained above, using the overall number of users in a municipality is preferable because it gives us much more data on user locations.

Next, in column (4), we winsorize the number of anti-refugee incidents at the 90th percentile to rule out the influence of outliers. This has no bearing on the results. In column (5), we implement the leave-one-out estimator typical for Bartik-type regressions by only using geo-tagged refugee posts *outside* of the municipality where we predict refugee attacks. The interaction term is highly significant with a  $t$ -stat of 3.27. In the last three columns, we replace the “Refugee posts” measure – which tracks refugee salience as posts on the AfD Facebook page – with alternative measures. Column (6) uses the number of comments on the AfD page containing the word refugee (“Refugee comments”). This picks up a slight but important difference, as the motivation for users to comment on existing posts might be different from creating new ones. In column (7), we weight the number of post by the number of likes they receive. Finally, in column (7) we use the *share* of refugee posts in all posts on the AfD page in a given week. The results are markedly similar independent of the used salience measure.

In Table A.19 in the online appendix, we next account in more detail for the skewed distribution our AfD Facebook use measure. As shown in Figure A.7, some municipalities have a large number of AfD users per capita. To rule out that these outliers are driving our results, we estimate our regression with several sample splits. We exclude municipalities with 0 users in column (1). We also show that our results hold for the sample below and above the median of AfD users with very similar estimated coefficients. They also hold when we exclude municipalities below the 10th and above the 90th percentile in AfD Facebook usage. Moreover, we show that the local social media exposure is monotonic, which is important since we are estimating the regressions using ordinary least squares. To do so, we divide the  $AfD\ Users/Pop.$  variable into four quartiles for municipalities with at least a single AfD user and interact them separately with the refugee salience trend. The first quartile serves as the excluded group. Strikingly, the results suggest that the effect of Germany-wide anti-refugee sentiment increases monotonically with the number of local Facebook users on the AfD page.

We also assess the robustness of the results to different transformations of the refugee attack variable and estimation methods in Table A.20. Note that we continue to include our baseline controls, which include an interaction term for population size. In column (1), we replicate our baseline finding using a dummy for all municipality-weeks with at least a single attack, estimated using OLS. In columns (2) and (3), we estimate the same model using logit regression with municipality fixed effects only and the Fernandez-Val and Weidner (2016) bias-adjusted estimator for two-way fixed effects. In a similar spirit, we use the log of one plus the number of attacks in column (4) to account for the left-skew in the ratio of incidents per asylum seeker. In all cases, the estimated coefficients are statistically significant; in fact, often more so than in the baseline regression.

In column (5), we more thoroughly address the concern that our findings may be influenced by the population scaling by again regressing  $\text{Log}(1 + \text{Refugee attacks})$  on the interaction of refugee posts and  $\text{Log}(1 + \text{AfD users})$ . This unscaled specification controls for population differences through the interaction of our salience measure with the number of inhabitants as a control variable. This regression yields an even stronger statistical association compared to our baseline result. A log-log version of the main equation Equation (1) yields similar results (column (6)).

Moreover, we analyze if our regression model is properly specified given the auto-correlation of the anti-refugee salience. In Table A.21 in the online appendix, we check whether leads or lags of anti-refugee salience have any effect after controlling for the effect of today's anti-refugee salience. To do so we use regress the residuals from our main regression on leads or lags of anti-refugee salience. Overall, none of the leads and lags has any significant effect over and above the effect of the contemporaneous salience.

In Table A.22, we also show a range of different standard errors for our baseline findings. As it turns out, clustering by municipality is an overall highly conservative choice. Finally, our results are robust to different levels of aggregation. In online appendix Table A.23 we show that the findings are similar if instead at the municipality level we aggregate our data at the county-level or up to the level of the 16 federal states.

### **3.8 How Many Refugee Attacks are Caused by Online Hate Speech?**

Our results allow us to provide a rough estimate of how many attacks against refugees would have taken place in a counter-factual world without right-wing social media posts. This calculation should only be interpreted as an illustration, rather than a precise exercise. Because we are interested in the impact of *right-wing* social media in particular, we use the  $\text{AfD users}/\text{Pop.}$  measure instead of the Nutella dummy. We leave a more quantitatively precise estimate for future research.

We calculate the predicted values of attacks if there were 0 anti-refugee posts on the AfD's Facebook page based on the conservative estimated coefficient of 0.151 in the regression with  $\text{county} \times \text{week}$  fixed effects. Multiplying with the coefficient with the  $\text{AfD users}/\text{Pop.}$  and  $\text{Refugee posts}$  gives us the estimated effect on the ratio of refugee attacks to asylum seekers. Again multiplying with the number of asylum seekers, then, is the predicted number of attacks due to right-wing social media salience. Our results imply that in the absence of anti-refugee posts on the AfD Facebook page 446 (13%) fewer anti-refugee incidents would have taken place. This back-of-the-envelope calculation further illustrates the non-negligible impact of social media.

## 4 Conclusion

Social media has quickly become a powerful tool for sharing and disseminating information online. In this paper, we investigate whether it can play a role in propagating violent hate crimes. Our findings suggest that social media has not only become a fertile soil for the spread of hateful ideas but also motivates real-life action. By combining detailed local data on Facebook usage with user-generated content, we can shed some light on how online posts are correlated with anti-refugee incidents in Germany. Plausibly exogenous variation in disruptions to users' internet or Facebook access – paired with a measure of social media penetration that is orthogonal to potential confounders – supports the view that some of this correlation reflects a causal effect.

Existing research shows that there is enormous persistence in local cultural attitudes towards foreigners (e.g. Becker and Pascali, 2016; Becker et al., 2016; Voigtlander and Voth, 2012, 2015). We extend this literature by showing that volatile, short-lived bursts in sentiment *within* a given location have substantial effects on people's behavior, and that social media may play a role in their propagation. Our findings are particularly timely in light of recent policy debates about whether and how to “regulate” hate speech on social media. Such legislation may come at a high price: since the lines between what constitutes free speech and hate speech are blurred, it can open the door for blanket censorship. Our work does, however, suggest that policymakers ignore online hate crime at their peril. It remains for future research to demonstrate effective ways to tackle online hate speech. By quantifying the extent of the problem, our paper aims to make a first step towards identifying potential harm arising from extended social media usage.

## References

- M. Adena, R. Enikolopov, M. Petrova, V. Santarosa, and E. Zhuravskaya. Radio and the rise of the nazis in prewar germany. *The Quarterly Journal of Economics*, 130(4):1885–1939, 2015.
- A. Alesina and E. La Ferrara. Ethnic diversity and economic performance. *Journal of Economic Literature*, 43:721–761, 2005.
- J. D. Angrist and J.-S. Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2008.
- P. Barberá. How social media reduces mass political polarization. evidence from germany, spain, and the us. *Job Market Paper, New York University*, 46, 2014.

- T. J. Bartik. Boon or boondoggle? the debate over state and local economic development policies. 1991.
- BBC. Social media warned to crack down on hate speech. 2017.
- S. O. Becker and L. Pascali. Religion, division of labor and conflict: Anti-semiticism in german regions over 600 years. *Religion*, (288), 2016.
- S. O. Becker, S. Pfaff, and J. Rubin. Causes and consequences of the protestant reformation. *Explorations in Economic History*, 62:1–25, 2016.
- M. Bhuller, T. Havnes, E. Leuven, and M. Mogstad. Broadband internet: An information superhighway to sex crime? *Review of Economic Studies*, 80(4):1237–1266, 2013.
- L. Boxell, M. Gentzkow, and J. M. Shapiro. Greater internet use is not associated with faster growth in political polarization among us demographic groups. *Proceedings of the National Academy of Sciences of the United States of America*, pages 10612–10617, 2017.
- D. Card and G. B. Dahl. Family violence and football: The effect of unexpected emotional cues on violent behavior. *The Quarterly Journal of Economics*, 126(1):103–143, 2011.
- T. Colussi, I. E. Isphording, and N. Pestel. Minority salience and political extremism. 2016.
- G. Dahl and S. DellaVigna. Does movie violence increase violent crime? *The Quarterly Journal of Economics*, pages 677–734, 2009.
- S. DellaVigna and M. Gentzkow. Persuasion: Empirical evidence. *Annual Review of Economics*, 2(1): 643–669, 2010.
- S. DellaVigna and E. La Ferrara. Economic and social impacts of the media, 2015.
- S. DellaVigna, R. Enikolopov, V. Mironova, M. Petrova, and E. Zhuravskaya. Cross-border media and nationalism: Evidence from serbian radio in croatia. *American Economic Journal: Applied Economics*, 6(3):103–32, July 2014.
- Destatis. *Statistisches Jahrbuch: Deutschland und Internationales*. Statistisches Bundesamt (Destatis), Wiesbaden, 2017. ISBN 978-3-8246-1057-0.

- R. Durante and E. Zhuravskaya. Attack when the world is not watching? us news and the israeli-palestinian conflict. *Journal of Political Economy*, Forthcoming, 2016.
- T. Eiseensee and D. Strmberg. News droughts, news floods, and u. s. disaster relief\*. *The Quarterly Journal of Economics*, 122(2):693–728, 2007.
- R. Enikolopov, A. Makarin, and M. Petrova. Social media and protest participation: Evidence from russia. 2016.
- I. Fernandez-Val and M. Weidner. Individual and time effects in nonlinear panel models with large n, t. *Journal of Econometrics*, 192(1):291 – 312, 2016.
- M. P. Fiorina and S. J. Abrams. Political polarization in the american public. *Annual Review of Political Science*, 11:563–588, 2008.
- V. Fouka and H.-J. Voth. Reprisals remembered: German-greek conflict and car sales during the euro crisis. 2013.
- A. Gavazza, M. Nardotto, and T. M. Valletti. Internet and politics: Evidence from uk local elections and local government policies. 2015.
- M. Gentzkow. Television and voter turnout. *The Quarterly Journal of Economics*, 121(3):931–972, 2006.
- R. Heintze. Social media atlas 2016/2017. *Faktenkontor GmbH*, 2017.
- S. Hölig and U. Hasebrink. *Reuters Institute Digital News Survey 2017: Ergebnisse für Deutschland*, volume Nr. 38 of *Arbeitspapiere des Hans-Bredow-Instituts*. Verlag Hans-Bredow-Institut, Hamburg, 2016. ISBN 978-3-87296-136-5.
- S. Hölig and U. Hasebrink. *Reuters Institute Digital News Survey 2017: Ergebnisse für Deutschland*, volume Nr. 42 of *Arbeitspapiere des Hans-Bredow-Instituts Nr. 42*. Verlag Hans-Bredow-Institut, Hamburg, 2017. ISBN 978-3-87296-142-6.
- S. Jha. Trade, institutions, and ethnic tolerance: Evidence from south asia. *American Political Science Review*, 107(4):806–832, 2013.
- M. Manacorda, A. Tesei, et al. Liberation technology: mobile phones and political mobilisation in africa. Technical report, Centre for Economic Performance, LSE, 2017.

- New York Times. How facebook warps our worlds, by Frank Bruni. 2016.
- New York Times. How fiction becomes fact on social media, by Benedict Carey. 2017a.
- New York Times. Seeking asylum in germany, and finding hatred, by Ainara Tiefenthaler, Shane Oneill and Andrew Michael Ellis. 2017b.
- A. Oksanen, J. Hawdon, E. Holkeri, M. Näsi, and P. Räsänen. Exposure to online hate among young social media users. *Soul of society: a focus on the lives of children & youth*, 18:253–273, 2014.
- Pew Research Center. Social media update 2016. Technical report, 2016.
- S. Stephens-Davidowitz. The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, 118:26–40, 2014.
- C. R. Sunstein. The law of group polarization. *Journal of Political Philosophy*, 10(2):175–195, 2002.
- C. R. Sunstein. *Republic. com 2.0*. Princeton University Press, 2009. ISBN 1400827833.
- C. R. Sunstein. *# Republic: Divided Democracy in the Age of Social Media*. Princeton University Press, 2017.
- The Guardian. Cps to crack down on social media hate crime, says Alison Saunders, by Vikram Dodd. 2017.
- N. Voigtlander and H.-J. Voth. Persecution perpetuated: The medieval origins of anti-semitic violence in nazi germany. *The Quarterly Journal of Economics*, 127(3):1339–1392, 2012.
- N. Voigtlander and H.-J. Voth. Nazi indoctrination and anti-semitic beliefs in germany. *Proceedings of the National Academy of Sciences of the United States of America*, 112(26):7931–7936, 2015.
- D. Yanagizawa-Drott. Propaganda and conflict: Evidence from the rwandan genocide \*. *The Quarterly Journal of Economics*, 129(4):1947–1994, 2014.

# **Online Appendix for “Fanning the Flames of Hate: Social Media and Hate Crime”**

## **A A Short History of the AfD**

The AfD was founded by Bernd Lucke, a professor of Economics at the University of Hamburg in 2013. Initially, the AfD positioned itself as an opposition party to the common European currency and the bailouts for Greece and Spain as a result of the financial crisis. Right from the start, however, the party also pandered to the right with a conservative social policy. Representatives of the AfD frequently attracted attention for using nationalist terminology and their attacks on the “Lügenpresse” (Lying Press) – a term popularized by the Nazis. With this political program and rhetoric, the AfD attracted 4.7% of the votes in the 2013 German Federal Election, only narrowly missing the 5% electoral threshold.

Nonetheless, the AfD celebrated several victories in state elections and winning seats in the state parliaments of Hesse, Saxony, Thuringia, Brandenburg, Bremen and Hamburg. Furthermore, the AfD reached 7.1% of the votes in the 2014 European Parliament election. With the cooling of the Euro Crisis, the focus of the party began to shift further to the right and topics like traditional family values or the role of Islam in Germany. These more nationalist-conservative political positions were mainly embodied by Frauke Petry and attracted a significant share of far-right recruits to the party. In 2015, Petry was eventually elected the main speaker of the party, which constituted a major defeat for its founder Bernd Lucke. As a result of the loss, Lucke resigned from his position as the leader of the party and left the party completely.

Several other important members of the AfD followed the example of its founder. The newly emerged leadership of the AfD mainly consisted of Frauke Petry, Alexander Gauland, Björn Höcke, and Beatrix von Storch, all of which occupy staunch national conservative positions. With the beginning of the refugee crisis, the AfD discussed ever more aggressively dangers of mass immigration and made clear their position that they were unwilling to accept any refugee into Germany. This messaging was accompanied by increased messages of xenophobia and criticism of Islam.

Under the new leadership and the impact of the refugee crisis, the AfD continued to win seats in local parliaments in several state elections in 2016 and secured seats in 14 of the 15 state parliaments. In the 2017 federal election, the AfD became the third strongest force in the German Parliament with 12.6% of the votes.

## **B Additional Details on the Data**

### **B.1. A Short Introduction to Facebook Pages and User Data**


On Facebook, celebrities, universities, restaurants, and political groups like the AfD have created their own Facebook pages. This page is the starting point for the followers of the AfD on Facebook. Any Facebook user who is interested in or supports the AfD can like the fan page. The messages posted on the AfD's page then will show up in the Facebook feed of that user. The Facebook feed consists of the individualized news and updates every user receives based on his friendship network and interest. In this way, the AfD is able to reach and rally their followers with political messages and party news.

In addition to being a means of communication for the AfD, the users can become active on the party's page as well. In general, there are three different ways for such interactions: First, people have the opportunity to post their own messages, links or pictures on the fan page. These posts are also visible for everybody, but they will not show up automatically in the Facebook feed of other users. Second, users can comment on posts and comments by other users or the AfD itself. Those comments show up beneath the original post and are also visible to the public. Third, each post or comment can be liked as a sign of support.

An example of how these three types of interactions show up on the AfD page can be found in Figure A.1. The Facebook Graph API allows to collect all post, comments, and likes from the AfD's fan page. In Figure A.1 the information collected using the Graph API are highlighted and labeled on the Facebook fan page. Facebook attributes each post to a unique user-id. This user-id makes it possible to attribute posts and comments to individual profiles.

For the hand collection of the user data, one has to visit each individual Facebook user profile. Depending on the users' privacy settings one can find among other information the place of residence and place of birth. Figure A.2 shows an example of a Facebook user profile and where to find the relevant information. If the user decided to hide this information, the box with the user information will be empty.

**Figure A.1: EXAMPLE OF ALTERNATIVE FOR GERMANY FACEBOOK PAGE**



*Notes:* This graphic shows an example of the Alternative for Germany’s Facebook fan page. The boxes and labels highlight the parts extracted using the Facebook Graph API. The names of users where removed by the authors to avoid privacy concerns.

**Figure A.2: EXAMPLE OF FACEBOOK USER PROFILE**



*Notes:* This graphic shows an example of a Facebook user profile. The box highlights the placement highlights the public user information extracted from Facebook. The personal information of users where removed by the authors to avoid privacy concerns.

**Table A.1: TRANSLATED EXAMPLE AFD POSTS FROM FACEBOOK**

<b>Date</b>	<b>Post</b>	<b># Likes</b>	<b># Comments</b>	<b># Shares</b>
19/05/2017	Side note in the local newspaper: A Turkish man (23) raped a young woman for more than four hours and was cleared of all charges by the judge. Verdicts that were only known in Arabic cultures are now finding their way into Germany. These pro multi cultural diversity judges are raping the German justice system, cultural sensitivity is apparently more important than the rule of law.	18917	302	1
10/05/2017	+++EUR 204.5 million per month for 500,000 asylum seekers paid in unemployment benefits+++ The top politicians of the old parties sold us the wave of migrants coming to Germany as an enrichment of culture and the economy, but the reality looks very different. The former skilled workers are being financed by social security because they cannot get a job because they are uneducated. Deportations are still not enforced and as a result everybody is fed by the state, even those without asylum.	2418	299	1446
27/12/2016	In Berlin, the police has arrested the wanted teenagers who are suspected to have set a homeless person on fire. Out of the seven suspects between the ages of 15 and 21, six are from Syria and one from Lybia. According to a report of the <i>Süddeutsche Zeitung</i> , all of them arrived in Germany as “refugees” between 2014 and 2016.	7984	1665	5725
15/11/2016	+ We knew it: More “refugees” are now coming via plane + The government currently has 500 migrants per month flown in from Italy and Greece. The Minister of Interior is further reviewing the admission of an additional 13,500 refugees from Turkey. Only a few European countries are complying with the EU directives, Germany - how could it be any different - is one of the first in line.	2153	1066	2584
21/10/2016	+++ Civil war in Garmisch-Partenkirchen? +++ The <i>Süddeutsche Zeitung</i> reports that the situation in Garmisch-Partenkirchen seems to be disastrous: “Blacks have taken over the power in the small skiing village in Germany”, the Kremlin-financed Russian station Russia Today reports. The French right-wing news portal Atlantic reports similar things about the alleged regime of dark-skinned refugees and the British Daily Mail speaks of riots in the streets , vandalism, and open sexual assaults.	2084	698	1926

*Notes:* This table reports 5 example posts from the AFD Facebook page that were posted by the party itself. The post were translated from German into English by the authors. The original posts can be found in Table A.2.


**Table A.2: EXAMPLE AFD POSTS FROM FACEBOOK**

<b>Date</b>	<b>Original Post</b>	<b># Likes</b>	<b># Comments</b>	<b># Shares</b>
19/05/2017	Randnotiz im Lokalteil: Über vier Stunden hinweg vergewaltigte ein Türke (23) eine junge Frau, von den Richtern erhielt er dafür einen Freispruch: Urteile, die man bisher nur aus dem arabischen Kulturkreis kannte, halten in Deutschland Einzug. Multi-Kulti-Richter vergewaltigen das deutsche Rechtssystem, Kultursensibilität steht scheinbar über der Staatsgewalt.	18917	302	1
10/05/2017	+++204,5 Millionen Euro* monatlich für 500.000 Asylbewerber in Hartz IV+++ Als Bereicherung für Kultur und Wirtschaft verkauften uns die Spitzenpolitiker der Altparteien die Migrationswelle nach Deutschland, die Realität sieht völlig anders aus. Die Fachkräfte von einst werden aus dem Sozialsystem finanziert, weil sie aufgrund ihres niedrigen Bildungsniveaus nicht in Arbeit gebracht werden können. Abschiebungen finden nach wie vor nicht statt, und so wird auch der staatlich versorgt, der gar kein Bleiberecht hat.	2418	299	1446
27/12/2016	In Berlin hat die Polizei die gesuchten Jugendlichen, die verdächtigt werden, einen Obdachlosen angezündet zu haben, vorläufig festgenommen. Sieben Tatverdächtige zwischen 15 und 21 Jahren sind es, sechs von ihnen stammen aus Syrien, einer aus Libyen. Einem Bericht der Süddeutschen Zeitung nach sind sie allesamt zwischen 2014 und 2016 als "Geflüchtete" nach Deutschland eingereist.	7984	1665	5725
15/11/2016	+ Also doch: Weitere "Flüchtlinge" kommen jetzt per Flugzeug + 500 Migranten pro Monat lässt die Bundesregierung im Moment aus Italien und Griechenland einfliegen. Der Innenminister prüft auch die Aufnahme von 13.500 Flüchtlingen aus der Türkei. Nur wenige europäische Länder fügen sich den EU-Beschlüssen, Deutschland ist - wie sollte es anders sein - wieder ganz vorne mit dabei.	2153	1066	2584
21/10/2016	+++ Bürgerkrieg in Garmisch-Partenkirchen? +++Die Süddeutsche Zeitung berichtet, dass es schlimm um Garmisch-Partenkirchen zu stehen scheint: "In dem kleinen Skiort in Deutschland hätten die Schwarzen die Macht übernommen, berichtet der kremlfinanzierte russische Auslandssender Russia Today auf seiner Internetseite. Das rechte französische Nachrichtenportal Atlantico weiß über das vermeintliche Regime dunkelhäutiger Flüchtlinge ähnliches zu berichten, und bei der britischen Daily Mail ist von Straenkämpfen, Vandalismus und sexuellen Übergriffen die Rede.	2084	698	1926


*Notes:* This table reports the 5 original German example posts from the AfD Facebook page that where posted by the AfD. The translated post can be found in Table A.1.

**Figure A.3: CORRELATION OF AfD AND NUTELLA USERS**


**(a) Correlation of AfD and Nutella Users/Pop.**




**(b) Raw Data**




**(c) Population below 150,000**



**(d) Population below 75,000**



**(e) Population below 25,000**



*Notes:* Panel (a) shows a bin scatter plot of the number of AfD users and Nutella users across the 4,466 municipalities in the sample, both scaled over population. Panels (b) through (e) plot the correlation of AfD and Nutella users (unscaled) at different points of the distribution in municipality population.

## B.2. Validation of Internet and Facebook Outages

Figure A.4: Map Internet Outages



*Notes:* This map plots the geographic distribution of internet outages per million inhabitants for the German municipalities in the data.

**Table A.3: Validation of Internet Outage Data**

**(a) Part 1/2**

<b>Date</b>	<b>Provider</b>	<b>Region</b>	<b>Description</b>	<b># Outages</b>	<b>Source</b>
12/06/2015	Kabel D. and Unitym.	Germany	The IT website “Netzwelt” reported about a large internet outage on June, 12th 2015. Users of the providers Kabel Deutschland and Unitymedia were especially affected. According to a spokesperson for Kabel Deutschland, the problem was caused by a disruption at the internet hub in Frankfurt.	61	Link
18/06/2015	Unitymedia	Karlsruhe	The news site “KA News” reported on July, 18th 2015 about a disruption at internet provider Kabel BW. Kabel BW, a subsidiary of Unitymedia, confirmed the problem and explained that their technicians were currently working to fix the problem. The outage affected the area of Karlsruhe.	36	Link
24/06/2015	Unitymedia	Cities in NRW	The “Rheinische Post” reported on June, 24th 2015 that many users of the provider Unitymedia encountered disruptions of the internet connection beginning on Wednesday, June 23th. Most of the reports came from the cities of Düsseldorf, Mönchengladbach, Neuss, and Münster. Unitymedia did not provide an official statement.	15	Link
05/07/2015	O2 and 1u1	Berlin	The IT website “Golem.de” reported on July, 5th 2015 that users of DSL provider O2 and 1und1 reported disruptions of their internet as well as phone connections. The problems had started on the previous Saturday and were largely fixed in the evening of the 5th. Both providers did not explain what had caused the problems.	27	Link
08/07/2015	Versatel	Münster	The “Haltemer Zeitung” reported on Juli 8th, 2015 that households in the city of Haltern were cutoff from the internet. The outage was caused by a damaged fiber optic cable. The same cable was also used by internet provider Unitymedia. As a result, Unitymedia users in the area of Münster were also affected by the problem.	29	Link
20/08/2015	Unitymedia	NRW and Hessen	The “Grißener Allgemeine” reported on August, 20th 2015 that many users of internet provider Unitymedia encountered disruptions of their internet connection beginning Wednesday, August 19th. The internet outage affected the entire state of Nordrhein-Westfalen as well as parts of Hessen. At the time of the report, Unitymedia was still investigating the cause of the outage.	81	Link

(b) Part 2/2

Date	Provider	Region	Description	# Outages	Source
04/12/2015	Telekom	Major cities	“ZDnet.de”, a website specialized in IT and electronics, reported on December 4th, 2015 that users of the internet provider German Telekom were encountering disruptions of their internet connections beginning in the early morning of the same day. Most of the reports came from the major cities Berlin, Hamburg, MÜNICH, and Frankfurt. According to the German Telekom, the problem was caused by a breakdown of a RADIUS server that is responsible for the authentication of internet access.	19	Link
30/06/2016	Vodafone and Kabel D.	Germany	“Heise.de”, the website from which we obtained the outage data, reported on June, 30th 2016 on an outage of the internet provider Kabel Deutschland that affected the entirety of Germany. The outage was caused by a problem at a computer cluster. At the time of the report the outage was still ongoing.	122	Link
21/07/2016	T-Online	Germany	The IT website “Golem.de” reported on July, 21th 2016 on internet problems of the provider T-Online. The outage did not only affect private households but also business customers of T-Online. A representative of the provider confirmed the problems but did not name any specific cause. At the time of the report technicians were still working to fix the problem.	41	Link
24/11/2016	O2	Germany	On November 24th, 2016 the website “Chip.de” reported that an Germany-wide outage of the internet provider O2 had occurred. The problems were concentrated in metropolitan areas. At the time of report O2 was still investigating the cause of the problem. Most likely cause of the outage was a problem with the VoIP-system of the company.	31	Link
27/11/2016	Telekom	Ruhr area	The “Spiegel” reported on November 27th, 2016 that users of the internet provider German Telekom were cut off from the internet. The outage mainly affected the Ruhr area, but internet problems were also reported in Frankfurt, Hannover, and Braunschweig. At the time of the report the Telekom was still working to correct the problem.	19	Link

*Notes:* This table reports several examples of internet outages that were reported in German newspapers or well-known specialized websites. Each entry lists the date of the outage as well as the affected provider and region. The table further features a short description of the outage and the link to the original news source. The column “# Outage” refers to the number of outages of the affected provider that users reported on Heise.de, which serve as the basis of our internet outage measures. The web pages were last visited in February 2018.

**Table A.4: Validation of Facebook Outage Data**

<b>Peak Date</b>	<b>Description</b>	<b>Source</b>
26/01/2015	The Facebook page was unavailable globally due to a server error. According to the official statement, the error “occurred after we introduced a change that affected our configuration systems.” Initially, the outage had been attributed to an attack by infamous hacker group “Lizard Squad”. The outage affected millions of users worldwide, including users of the Facebook messenger, Instagram, and the dating app Tinder (which uses Facebook data).	Link
27/03/2015	Facebook displayed an error message that the site “is down for required maintenance right now”, likely the result of a service disruption. The outage was concentrated in Western Europe, particularly in Germany, the Netherlands, and the United Kingdom.	Link
15/07/2015	Facebook suffered another worldwide outage, showing users a simple “Service Unavailable” message. The outage affected all services including the popular Facebook messenger. Although the initial issue was resolved relatively quickly, the problems persisted for many users.	Link
29/09/2015	Users experienced extremely slow or no access to Facebook after a previous disruption on September 24. User reports and the news coverage indicate that Germany was particularly badly hit. In a statement to CNBC, Facebook acknowledged the outage and explained that “configuration problems” were at the root of it.	Link,Link2
14/03/2016	Users in Western Europe - particularly Germany, Austria, Poland, the Netherlands, Belgium, and the United Kingdom - were barred from logging into or commenting on Facebook. The Facebook app was apparently particularly affected.	Link
16/06/2016	Facebook outage concentrated in Western Europe. Users were unable to log in, post, or use the messenger and could not access Facebook pages (including that of the AfD).	Link
14/09/2016	Worldwide Facebook outage, affecting almost the entire European continent and the eastern United States. Users were unable to log in or post and read content.	Link,Link2
13/01/2017	Users in Western Europe and the eastern United States experienced widespread issues in accessing Facebook, particularly from computer devices.	Link

*Notes:* This table lists the dates of the major Facebook outages that occurred during our sample period. The links lead to the news articles used to identify the disruptions.


### B.3. Additional Variable Overview

**Table A.5: SUMMARY STATISTICS FOR ADDITIONAL CONTROLS**

	Level	Obs.	Mean	S.D.	Min.	Max.
<b>Additional Media and Internet Controls<sup>†</sup></b>						
Internet outages/Pop.	Muni.-Week	495,726	0.071	0.707	0	36.138
Registered domains/Pop.	County	495,726	12.741	5.013	5.142	125.226
News paper sales/Pop.	Municipality	368,520	870.943	874.001	0	16442.730
<b>Additional Right-Wing Controls</b>						
Nazi murders/Pop. <sup>†</sup>	Municipality	495,726	0.015	0.121	0	2.822
NPD vote share (2017)	Election Distr.	495,726	0.491	0.410	0	2.006
Log(1+Deported Jews)	Municipality	495,726	0.606	1.350	0	10.930
Log(1+Stürmer letters)	Municipality	495,726	0.125	0.449	0	5.872
<b>Additional Socio-Economic Controls</b>						
Average age	Municipality	479,853	44.971	2.277	27	56.200
Share benefit recipients	Municipality	495,726	0.382	0.168	0.0462	1.087
Share non-Christians	Municipality	479,853	46.219	2.522	26.8	57.700
Manufacturing share	County	493,062	26.913	9.315	2.1431	57.668
<b>Additional Voting Controls (2017 Election)</b>						
CDU vote share	Election Distr.	495,726	35.191	5.510	22.152	52.729
SPD vote share	Election Distr.	495,726	18.501	6.106	7.7011	37.395
Left vote share	Election Distr.	495,726	8.297	4.312	4.1354	20.854
Green vote share	Election Distr.	495,726	7.514	3.310	2.1961	21.013
FDP vote share	Election Distr.	495,726	9.788	2.413	5.4062	17.444
Pirate vote share	Election Distr.	495,726	0.322	0.151	0	0.803
Voter turnout	Election Distr.	495,726	76.445	3.136	65.929	83.881
<b>Additional Demographic Controls</b>						
Share aged 3-14	Municipality	495,726	11.111	2.616	0	18.207
Share aged 15-17	Municipality	495,726	3.073	0.944	0	6.953
Share aged 18-24	Municipality	495,726	7.475	1.748	0	16.799
Share aged 25-29	Municipality	495,726	4.971	1.237	0	12.169
Share aged 30-39	Municipality	495,726	10.647	2.224	0	19.192
Share aged 40-49	Municipality	495,726	16.668	3.298	0	24.373
Share aged 50-64	Municipality	495,726	20.667	4.200	0	31.469
Share above 65	Municipality	495,726	19.903	4.897	0	38.127

*Notes:* This table reports summary statistics for the additional control variables in the estimation sample. Variables tagged with a † are scaled by population in 10,000. Share variables are in percent.

**Figure A.5: DAILY INTERNET USERS AND SHARE OF HOUSEHOLDS WITH BROADBAND ACCESS**



*Notes:* This figure plots the share of households with access to broadband internet ( $\geq 16$  Mbit/s) against the percentage of individuals using the internet daily across from Eurostat survey data, binned into 16 quantiles. The corresponding correlation coefficient is 0.9245.

**Table A.6: Overview Variables**

(a) Part 1/4

Variable	Level	Description	Source
<b>Refugee Attacks</b>			
Refugee Attacks/Refugees	Muni.-Week	Constructed by dividing the number of anti-refugee incident in a municipality and week by the number of refugees.	Amadeu Antonio Foundation
Arson Attacks/Refugees	Muni.-Week	Same as Refugee Attacks/Refugees but limited to arson attacks as classified by the Amadeu Antonio Foundation.	Amadeu Antonio Foundation
Other Property Attack/Refugees	Muni.-Week	Same as Refugee Attacks/Refugees but limited to attacks leading to miscellaneous property damages as classified by the Amadeu Antonio Foundation.	Amadeu Antonio Foundation
Assaults/Refugees	Muni.-Week	Same as Refugee Attacks/Refugees but limited to assault as classified by the Amadeu Antonio Foundation.	Amadeu Antonio Foundation
Demonstrations/Refugees	Muni.-Week	Same as Refugee Attacks/Refugees but limited to demonstrations as classified by the Amadeu Antonio Foundation.	Amadeu Antonio Foundation
Suspected Cases/Refugees	Muni.-Week	Same as Refugee Attacks/Refugees but limited to suspected attacks still under investigation, as classified by the Amadeu Antonio Foundation.	Amadeu Antonio Foundation
<b>Social Media Data</b>			
AfD Users/Pop.	Municipality	The number of AfD Users in each municipality divided by population.	Facebook
Nutella Users/Pop.	Municipality	The number of Nutella Users in each municipality divided by population.	Facebook
$I_{Many\ Nutella\ Users}$	Municipality	The <i>Many Nutella Users Dummy</i> is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise.	Facebook
Refugee Posts	Week	The number of posts on the AfD Facebook page that contain the word 'Flüchtling' (refugee) in a given week.	Facebook
Posts/AfD Users	Municipality	The total number of posts attributed to AfD users of a given municipality divided by the number of AfD Facebook users.	Facebook
Comments/AfD Users	Municipality	Total number of comments that posts by AfD users of a given municipality received divided by the number of AfD Facebook users.	Facebook
Likes/AfD Users	Municipality	The total number of likes that posts by AfD users in a given municipality received divided by the number of AfD Facebook users.	Facebook

(b) Part 2/4

Variable	Level	Description	Source
<b>Auxiliary Variables</b>			
$I_{Internet\ Outage}$	Muni.-Week	Dummy variable that is equal to 1 for municipality-week observations that are in the top quartile of the reported internet outages per capita ratio, and 0 otherwise. The number of user-reported outages comes from Heise.de. We exclude outages that are shorter than 24 hours.	Heise.de
$I_{Facebook\ Outage}$	Week	Dummy variable that is equal to 1 for the weeks with major Facebook outages as described in Table A.4, and 0 otherwise.	Various news sources www.durchgezaehlt.org
PEGIDA Demonstrators	Week	The number of demonstrators at the weekly PEGIDA (Patriotic European Against the Islamisation of the West) demonstrations in Dresden.	Newspaper websites
Refugee news coverage	Week	First principal component of the number of news items about refugees on the websites of the major newspapers Frankfurter Allgemeine Zeitung, Der Spiegel, Zeit, Tagesschau, and Handelsblatt.	Compact Magazin website
Compact news coverage	Week	The number of news items about refugees on the website the right-wing Compact Magazin.	Compact Magazin website
<b>Baseline Controls</b>			
Population	Municipality	The population of each municipality in 2015 from the shape file of the BKG. The population numbers in the shape file are equivalent to the 2015 data from the German Statistical Office (Destatis).	BKG/Destatis
GDP/Worker	County	GDP per working population at the county-level.	Destatis
Population density	Municipality	Population density, defined as population over municipality size (in $km^2$ ).	Destatis
AfD vote share (2017)	Election Distr.	The share of votes cast for the AfD in the 2017 German Federal Parliament Election.	Destatis
Share Abitur	Municipality	The share of people whose highest educational attainment is at least "Abitur", the German high-school certificate.	Destatis
Share Broadband access	County	The share of the population that have access to at least 16 Mbit/s internet connection speed.	BMVI, TÜV Rheinland
Share immigrants	Municipality	The share of the population that are immigrants.	Destatis
<b>Raw Data</b>			
Refugee attacks	Muni.-Week	The number of anti-refugee incident in a municipality and week.	Amadeu Antonio Foundation
Population (2015)	Municipality	The population for each municipality.	BKG
Refugees (2015)	County	The number of asylum seekers in each county.	Destatis
AfD Users	Municipality	The number of users of the AfD Facebook page we could locate based on their reported place of residence.	Facebook
Nutella Users	Municipality	The number of users of the Nutella Facebook page we could locate based on their reported place of residence.	Facebook

## (c) Part 3/4

Variable	Level	Description	Source
<b>Additional Media and Internet Controls</b>			
Internet outages/Pop.	Municipality	The total number of large internet outages (as defined above) that occurred in a municipality over the sample period	Heise.de
Registered domains/Pop.	County	The number of registered .de domains in a given county, divided by the county population.	Destatis
Newspaper sales/Pop.	Municipality	The number of newspaper copies sold in a given municipality, divided by the population. The data do not contain information for municipalities smaller than 3000 inhabitants.	ZMG
<b>Additional Right-Wing Controls</b>			
Nazi murders (1990-2016)	Municipality	The number of murders classified as having a neo-Nazi motive in a municipality between 1990 and 2016, scaled by population.	Mut gegen rechte Gewalt
NPD vote share (2017)	Election Distr.	The share of votes cast for the extremist right-wing NPD (National Democratic Party of Germany) in the 2017 German Federal Parliament Election.	Bundeswahlleiter
Log(1+Deported Jews)	Municipality	The natural logarithm of the number of Jews that were deported during Nazi times. For the analysis of cross-sectional correlates, we scale the number of deported Jews by the number of Jews in a municipality in 1933 (see text for details).	Voigtlander and Voth (2012)
Log(1+Stürmer letters)	Municipality	The natural logarithm of the number of letters written to "Der Stürmer", the antisemitic newspaper published by Nazi politician Julius Streicher. For the analysis of cross-sectional correlates, we scale the number of letters by the population in 1933 (see text for details).	Voigtlander and Voth (2012)
<b>Additional Socio-Economic Controls</b>			
Average age	Municipality	The average age in each municipality.	Destatis
Benefit recipients/Pop.	Municipality	The number of social benefit recipients in a given municipality divided by the population.	Destatis
Non-christians/Pop.	Municipality	The number of non-Christians in a given municipality divided by population.	Destatis
Manufacturing share (%)	County	The share of employees in manufacturing in a given county.	Destatis

(d) Part 4/4

Variable	Level	Description	Source
<b>Additional Voting Controls (2013 &amp; 2017 Election)</b>			
CDU vote share	Election Distr.	The share of votes cast for the CDU in the 2017 German Federal Parliament Election.	Bundeswahlleiter
SPD vote share	Election Distr.	The share of votes cast for the SPD in the 2017 German Federal Parliament Election.	Bundeswahlleiter
Left vote share	Election Distr.	The share of votes cast for the party "Die Linke" (The Left) in the 2017 German Federal Parliament Election.	Bundeswahlleiter
Green vote share	Election Distr.	The share of votes cast for the party "B90/Die Grünen" (Green Party) in the 2017 German Federal Parliament Election.	Bundeswahlleiter
FDP vote share	Election Distr.	The share of votes cast for the FDP in the 2017 German Federal Parliament Election.	Bundeswahlleiter
Pirate vote share	Election Distr.	The share of votes cast for the Pirate party in the 2017 German Federal Parliament Election.	Bundeswahlleiter
Voter turnout	Election Distr.	The voter turnout in the 2017 German Federal Parliament Election.	Bundeswahlleiter
<b>Additional Demographic Controls</b>			
Share aged 3-14	Municipality	The number of people in the 3-14 age group, divided by population.	Destatis
Share aged 15-17	Municipality	The number of people in the 15-17 age group, divided by population.	Destatis
Share aged 18-24	Municipality	The number of people in the 18-24 age group, divided by population.	Destatis
Share aged 25-29	Municipality	The number of people in the 25-29 age group, divided by population.	Destatis
Share aged 30-39	Municipality	The number of people in the 30-39 age group, divided by population.	Destatis
Share aged 40-49	Municipality	The number of people in the 40-49 age group, divided by population.	Destatis
Share aged 50-64	Municipality	The number of people in the 50-64 age group, divided by population.	Destatis
Share above 65	Municipality	The number of people in the 65+ age group, divided by population.	Destatis

## B.4. Additional Results for Balancedness

**Table A.7: BALANCEDNESS OF SOCIAL MEDIA USAGE - MEDIAN SAMPLE SPLITS**

Above/below median in:	AfD Users/Pop.			Nutella Users/Pop.			$I_{Many\ Nutella\ Users}$		
	Below	Above	<i>t</i> -stat	Below	Above	<i>t</i> -stat	No	Yes	<i>t</i> -stat
<b>Baseline Controls</b>									
Population (2015) †	10326.509	26474.066	-7.256***	9678.282	27122.293	-7.846***	11568.081	27930.720	-6.148***
GDP/Worker	63636.758	62552.057	3.678***	63345.89	62843.543	1.701*	62899.839	63370.928	-1.574
Population density	231.655	332.186	-8.877***	227.665	336.176	-9.596***	244.547	333.651	-7.311***
AfD vote share (2017)	13.615	14.816	-6.735***	14.728	13.703	5.741***	14.332	14.053	1.551
Share Abitur	28.066	30.011	-7.929***	27.871	30.205	-9.546***	29.153	28.862	1.145
Share broadband access	83.185	82.814	1.161	82.396	83.603	-3.788***	82.777	83.301	-1.620
Share immigrants	13.995	13.930	0.283	13.177	14.743	-6.809***	13.810	14.171	-1.529
<b>Additional Media and Internet Controls†</b>									
Internet outages/Pop.	0.082	0.06	1.02	0.065	0.077	-0.546	0.061	0.085	-0.981
Registered domains/Pop.	12.83	12.653	1.184	12.728	12.754	-0.174	12.64	12.876	-1.511
News paper sales/Pop.	978.601	771.157	6.837***	991.291	769.93	7.343***	919.101	808.166	3.581***
<b>Additional Right-Wing Controls</b>									
Nazi murders/Pop.†	0.01	0.021	-3.124***	0.014	0.017	-0.887	0.014	0.017	-0.993
NPD vote share (2017)	0.446	0.536	-7.349***	0.501	0.482	1.57	0.502	0.476	2.133**
Deported Jews/Jews (1933)	38.931	31.351	1.699*	34.594	33.685	0.234	34.011	34.049	-0.011
Stürmer letters/Pop. (1933)	2.347	1.497	2.553**	2.075	1.664	1.317	1.89	1.730	0.557
<b>Additional Socio-Economic Controls</b>									
Average age	44.655	45.292	-9.288***	44.904	45.039	-1.955*	45.023	44.900	1.757*
Share benefit recipients	0.004	0.004	-5.068***	0.004	0.004	-6.042***	0.004	0.004	0.591
Share non-Christians	45.818	46.626	-10.669***	46.085	46.354	-3.509***	46.248	46.181	0.864
Manufacturing share (%)	27.698	26.128	5.634***	27.28	26.546	2.626***	27.033	26.743	1.021
<b>Additional Voting Controls (2017 Election)</b>									
CDU vote share	35.992	34.389	9.825***	35.557	34.824	4.454***	35.134	35.269	-0.804
SPD vote share	18.376	18.625	-1.364	17.417	19.585	-12.05***	18.495	18.511	-0.086
Left vote share	7.727	8.867	-8.906***	8.269	8.325	-0.429	8.346	8.229	0.892
Green vote share	7.883	7.146	7.482***	7.531	7.498	0.334	7.465	7.582	-1.158
FDP vote share	9.878	9.698	2.499**	9.683	9.892	-2.896***	9.795	9.778	0.238
Pirate vote share	0.322	0.323	-0.172	0.312	0.332	-4.47***	0.320	0.326	-1.364
Voter turnout	76.87	76.019	9.155***	76.614	76.275	3.626***	76.44	76.449	-0.093
<b>Additional Demographic Controls</b>									
Share aged 3-14	11.528	10.695	10.785***	11.197	11.026	2.185**	11.13	11.086	0.564
Share aged 15-17	3.248	2.898	12.629***	3.092	3.055	1.308	3.082	3.061	0.760
Share aged 18-24	7.555	7.394	3.075***	7.388	7.562	-3.339***	7.369	7.621	-4.784***
Share aged 25-29	4.893	5.048	-4.185***	4.854	5.087	-6.325***	4.843	5.148	-8.039***
Share aged 30-39	10.744	10.550	2.911***	10.638	10.656	-0.272	10.565	10.760	-2.919***
Share aged 40-49	17.000	16.336	6.763***	16.787	16.549	2.409**	16.726	16.588	1.397
Share aged 50-64	20.671	20.664	0.059	20.69	20.644	0.362	20.701	20.621	0.638
Share above 65	19.442	20.364	-6.321***	19.61	20.196	-4.006***	19.816	20.025	-1.422

*Notes:* This table reports the arithmetic mean of the variables on the left column, where municipalities are split into above and below median values of the ratio of AfD and Nutella users to population or the many Nutella users dummy. In column “*t*-stat”, we report the *t*-statistic of a test for equal means under the null hypothesis that the average in the groups is equivalent. See the data section for details on variable construction.

## C Additional Results: AfD Activity and Anti-Refugee Incidents

**Table A.8: SOCIAL MEDIA AND HATE CRIME: TIME SERIES CORRELATIONS**

Dependent variable:	Number Refugee attacks		
	All municipalities (1)	AfD users > 0 (2)	AfD users = 0 (3)
Refugee posts	94.61*** (14.88)	94.78*** (14.70)	71.64*** (19.18)
Observations	111	111	111
Adjusted $R^2$	0.335	0.336	0.188

*Notes:* This table reports the results of a time series regression of the number of refugee attacks on the number of anti-refugee posts on the AfD's Facebook page. We standardize the number of refugee attacks to make the coefficients comparable across the different samples. The regression is reported for the full sample, for municipalities with at least one AfD Facebook user and for municipalities with no AfD Facebook users. Robust standard errors are in parentheses. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.9: CROSS-SECTIONAL CORRELATIONS: AfD AND NUTELLA USERS**

	AfD Users/Pop.	
	(1)	(2)
Nutella Users/Pop.	0.459*** (0.117)	
$I_{Many\ Nutella\ Users}$		0.882*** (0.091)
Observations	4,466	4,465
Adj. $R^2$	0.107	0.023
Kleibergen-Paap F stat.	15.472	95.084

*Notes:* This table presents results of OLS regressions of the ratio of AfD users to population on Nutella users to population. The *Many Nutella Users Dummy* is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise. Robust standard errors are in parentheses. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.10: MAIN RESULTS WITH PLOTTED CONTROLS**

Dependent variable Social media measure	Refugee Attacks/Asylum Seekers		
	AfD Users/Pop. (1)	Nutella Users/Pop. (2)	$I_{Many\ Nutella\ Users}$ (3)
Social media users/Pop. $\times$ Refugee posts	0.369*** (0.113)	0.287*** (0.091)	2.317*** (0.621)
Population $\times$ Posts	0.002 (0.001)	0.002 (0.001)	0.001 (0.001)
GDP/Worker $\times$ Posts	-0.010** (0.005)	-0.010** (0.005)	-0.011** (0.005)
Pop. density $\times$ Posts	0.426*** (0.132)	0.443*** (0.133)	0.426*** (0.129)
AfD vote share $\times$ Posts	0.953 (0.752)	1.237 (0.784)	1.207 (0.784)
Share Abitur $\times$ Posts	0.761* (0.427)	0.761* (0.429)	0.922** (0.444)
Share broadband access $\times$ Posts	0.012** (0.006)	0.012** (0.006)	0.012** (0.006)
Share immigrants $\times$ Posts	-1.096 (0.668)	-1.130* (0.669)	-1.056 (0.667)
Week FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Observations	480,963	480,963	480,852
Number of municipalities	4,322	4,322	4,321
$R^2$	0.045	0.045	0.045

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* and *Nutella users/Pop.* are the ratio of people with any activity on the Facebook pages of the AfD and Nutella, respectively, scaled over the total municipality population. The *Many Nutella Users Dummy* is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. The interaction term coefficients for population, GDP/worker, and population density are multiplied by 1 million; the other control variable interactions by 1,000. Standard errors are calculated as indicated. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.11: SOCIAL MEDIA EFFECTS: VIOLENT VS. NON-VIOLENT INCIDENTS**

	Type of Incident				
	Arson (1)	Property damage (2)	Assault (3)	Demonstration (4)	Suspected cases (5)
AfD users/Pop. $\times$ Refugee posts	0.027** (0.012)	0.227*** (0.072)	0.078** (0.032)	0.039 (0.029)	-0.001 (0.002)
Week FE	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts	Yes	Yes	Yes	Yes	Yes
Observations	480,963	480,963	480,963	480,963	480,963
Number of municipalities	4,333	4,333	4,333	4,333	4,333
$R^2$	0.012	0.041	0.023	0.017	0.009

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on local social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. See text for an explanation of the control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.


**Table A.12: REPLICATION WITH OTHER SOCIAL MEDIA PROXIES AND ADDITIONAL CONTROL VARIABLES**

	Additional interacted controls								
	Baseline controls only (1)	Right-wing controls (2)	Media controls (3)	Socio-economic controls (4)	2017 vote controls (5)	Age structure controls (6)	All controls (7)	County × Week FE (8)	County × Week FE† (9)
<b>Panel A: Nutella Users/Pop.</b>									
Nutella users/Pop. × Refugee posts	0.287*** (0.091)	0.283*** (0.092)	0.246* (0.126)	0.258*** (0.093)	0.263*** (0.091)	0.129 (0.082)	0.113 (0.120)	0.119* (0.066)	–
<b>Panel B: <math>I_{Many\ Nutella\ Users}</math></b>									
$I_{Many\ Nutella\ Users} \times$ Refugee posts	2.317*** (0.621)	2.297*** (0.617)	2.213*** (0.681)	2.228*** (0.604)	2.182*** (0.600)	1.552*** (0.527)	1.381** (0.604)	0.373 (0.311)	0.947*** (0.346)
Week FE		Yes	Yes	Yes	Yes	Yes	Yes	–	–
Municipality FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) × Posts	Yes								
Right-wing controls (4) × Posts		Yes							
Media controls (3) × Posts			Yes						
Socio-econ. controls (4) × Posts				Yes					
Election controls (7) × Posts					Yes				
Age controls (8) × Posts						Yes			
County × Week FE								Yes	Yes


*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on local social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *Nutella users/Pop.* are the ratio of people with any activity on the Facebook page of Nutella, scaled over the total municipality population. The *Many Nutella Users Dummy* is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. See text and table Table 2 for the included control variables. The † in column (9) indicates that we replace  $I_{Many\ Nutella\ Users}$  with a dummy that is 1 for the top tercile of Nutella users *within a county*, 0 for the bottom tercile, and missing for the middle tercile. This is to create additional within-county variation to counter the *county × week* fixed effects: the within-county standard deviation of the baseline dummy is 0.12, compared to 0.27 with the “adjusted” dummy that takes out the middle tercile. All results in the paper are robust to using this adjusted measure (available upon request). Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Figure A.6: Randomization Test for Internet Outages**

**(a) Within Municipalities with internet Outages**



**(b) All Municipalities**



*Notes:* This figure shows the results of the randomization test, where we randomly assign placebo internet outages to 313 municipality-week pairs. We repeat this process 500 times and save  $t$ -stat of the triple interaction term of interest. Panel (a) shows a histogram of the  $t$ -stats for the municipalities in which a large internet outage occurred in our data. Panel (b) show the histogram of the  $t$ -stats for the entire sample of municipalities. The vertical line marks a  $t$ -stat of -1.96.

**Table A.13: DO INTERNET OUTAGES REDUCE FACEBOOK ACTIVITY? MORE EVIDENCE**

	Refugee posts		Total posts		Total likes		Total comments		Total shares	
	Full	Posts > 0	Full	Posts > 0	Full	Likes > 0	Full	Comments > 0	Full	Shares > 0
<b>Panel A: Outages Above Median</b>										
No disruption	0.008	0.144	0.069	0.283	0.153	0.860	0.071	0.370	0.157	2.219
Disruption	0.002	0.019	0.016	0.034	0.055	0.155	0.026	0.068	0.002	0.011
Observations	495,726	27,306	495,726	120,435	495,726	88,356	495,726	94,572	495,726	35,187
Diff. > 0?	3.824***	6.337***	9.426***	17.147***	3.009***	7.261***	3.957***	9.543***	4.602***	4.632***
<b>Panel B: Outage in Top Quartile (Used in Regression)</b>										
No disruption	0.008	0.144	0.069	0.282	0.153	0.859	0.071	0.370	0.157	2.214
Disruption	0.003	0.111	0.003	0.013	0.000	0.000	0.000	0.000	0.000	0.000
Observations	495,726	27,306	495,726	120,435	495,726	88,356	495,726	94,572	495,726	35,187
Diff. > 0?	1.466*	0.291	16.397***	16.845***	26.264***	26.349***	29.272***	29.380***	4.654***	4.656***

*Notes:* This table presents the arithmetic mean of local measures of Facebook activity based on linking user locations with posts, likes, comments, and shares. The row “Diff. > 0” presents the results of one-sided t-tests under the null hypothesis that the mean values are larger in the “No outage” sample, and the associated p-values (we do not assume equal variances across samples). The samples indicated as “> 0” are restricted to municipalities where we observe at least one post, like, comment, or share over time, respectively.

**Table A.14: ROBUSTNESS: ALTERNATIVE DEFINITIONS OF INTERNET OUTAGES**

	AfD users/Pop. (1)	Nutella users/Pop. (2)	$I_{Many\ Nutella\ Users}$ (3)
Panel A: Continuous measure of internet outages/Pop. (1,249 Outages)			
Social media users/Pop. $\times$ Refugee posts	0.369*** (0.113)	0.288*** (0.091)	2.325*** (0.622)
Social media users/Pop. $\times$ Posts $\times$ Outage	-0.019* (0.011)	-0.020* (0.012)	-0.055 (0.041)
Panel B: Outage in top quartile (baseline) (313 Outages)			
Social media users/Pop. $\times$ Refugee posts	0.533*** (0.151)	0.428*** (0.119)	2.889*** (0.675)
Social media users/Pop. $\times$ Posts $\times$ Outage	-1.917** (0.804)	-1.330** (0.622)	-6.966*** (2.512)
Panel C: Outage in top quartile including shorter Outages (597 Outages)			
Social media users/Pop. $\times$ Refugee posts	0.369*** (0.113)	0.287*** (0.091)	2.322*** (0.622)
Social media users/Pop. $\times$ Posts $\times$ Outage	-1.597** (0.679)	-0.834* (0.466)	-3.544** (1.802)
Panel D: Outage in top 5% (alternative cutoff) (63 Outages)			
Social media users/Pop. $\times$ Refugee posts	0.369*** (0.113)	0.287*** (0.091)	2.318*** (0.621)
Social media users/Pop. $\times$ Posts $\times$ Outage	-1.641*** (0.487)	-0.563 (0.722)	-5.719* (3.236)
Week FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Baseline controls $\times$ Posts	Yes	Yes	Yes

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media measures, where we use exogenous variation in internet outages. The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* and *Nutella users/Pop.* are the ratio of people with any activity on the respective Facebook pages to population. The *Many Nutella Users Dummy* is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. In panel A, we use the number of user-reported internet disruptions in a week scaled over municipality population as a measure of outage. In panel B, we use the baseline dummy explained above, i.e. outages in the top quartile. In Panel C, we include outages shorter than 24 hours (as discussed in Section 2 we exclude this for our baseline measures) and define a new dummy for outages in the top quartile. In Panel D, we only use the top 5% of outages (scaled to population) as a dummy variable. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.15: CORRELATION OF ALTERNATIVE SENTIMENT MEASURES**

	Refugee posts	Refugee news	Compact articles
Refugee posts	1		
Refugee news	0.7919*	1	
Compact articles	0.7139*	0.6287*	1
PEGIDA	-0.1029	0.0227	-0.0825

*Notes:* This table presents the Pearson correlation coefficients for the alternative refugee salience measures discussed in section 3.6. *Refugee posts* is our baseline measure based on the AfD Facebook page. *Refugee news* is the first principal component of news coverage described in the data section. *Compact articles* is the number of articles about refugees on the right-wing publication “Compact”. *PEGIDA* is the weekly number of protest participants of the “Patriotic Europeans Against the Islamisation of the West” movement. \* indicates statistical significance at the 0.01 level.

**Table A.16: EXPLORING ALTERNATIVE SENTIMENT MEASURES – HORSE RACE**

	Baseline (1)	Media coverage (2)	Right-wing media coverage (3)	Right-wing PEGIDA sentiment (4)	All measures (5)
AfD users/Pop. × Refugee posts	0.369*** [2.264] (0.113)	0.273* [1.676] (0.154)	0.304*** [1.867] (0.109)	0.398*** [2.446] (0.122)	0.332* [2.038] (0.193)
AfD users/Pop. × Media coverage		0.387 [0.742] (0.409)			0.104 [0.199] (0.442)
AfD users/Pop. × Compact articles			0.188 [0.556] (0.173)		0.100 [0.297] (0.184)
AfD users/Pop. × PEGIDA demonstrators				0.450** [1.770] (0.193)	0.445** [1.749] (0.196)
Week FE	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes
Baseline controls × Posts	Yes	Yes	Yes	Yes	Yes
Observations	480,963	480,963	480,963	381,304	381,304
Number of municipalities	4,333	4,333	4,333	4,333	4,333
$R^2$	0.045	0.045	0.045	0.054	0.054

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers; *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. *Pegida* is the weekly number of “PEGIDA” demonstrators in Dresden; the lower observation count stems from a few weeks with missing data. *Media coverage* is the first principal component of the number of news items on major German newspaper websites containing the word refugee (“Flüchtling”). *Compact articles* is the equivalent measure based on the right-wing Compact Magazin’s website. See text for an explanation of the control variables. Standardized coefficients reported in square brackets, based on variable transformations with a mean of 0 and a standard deviation of 1. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.17: EXPLORING INDIVIDUAL NEWS COVERAGE MEASURES (ONE-BY-ONE)**

		General Media Coverage Measure						
Principal Component (Baseline)		FAZ Only	Spiegel Online Only	Zeit Only	Tagesschau Only	Handelsblatt Only	Include All	
		(2)	(3)	(4)	(5)	(6)	(7)	
AfD users/Pop. × Refugee posts	0.408** (0.173)	0.425*** (0.158)	0.366** (0.162)	0.400** (0.160)	0.575*** (0.169)	0.444*** (0.150)	0.375** (0.161)	
AfD users/Pop. × Media coverage	0.507 (0.408)							
AfD users/Pop. × FAZ articles		0.145 (0.117)					-0.0199 (0.119)	
AfD users/Pop. × Spiegel Online articles			0.0343* (0.0203)				0.0488 (0.0311)	
AfD users/Pop. × Zeit articles				0.0289 (0.0217)			-0.0146 (0.0313)	
AfD users/Pop. × Tagesschau articles					-0.361 (0.281)		-0.875** (0.369)	
AfD users/Pop. × Handelsblatt articles						0.202 (0.127)	0.236* (0.130)	
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Observations	495,726	495,726	495,726	495,726	495,726	495,726	495,726	
Number of municipalities	4,333	4,333	4,333	4,333	4,333	4,333	4,333	
R <sup>2</sup>	0.045	0.045	0.045	0.045	0.045	0.045	0.045	


*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers; *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. *Media coverage* is the first principal component of the number of news items on all German newspaper websites included containing the word refugee (“Flüchtling”). The other variables control for the number of news items individually by outlet. See text for an explanation of the control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.18: FURTHER ROBUSTNESS CHECKS**

	Refugee Attacks/ Asylum Seekers	Pop.- Weighted Regression	Pre- Sample Users	Winsorize Attacks	Leave One Out Estimator	Refugee Comments Interaction	Refugee Likes Interaction	Refugee Post Share
Lagged Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AfD users/Pop. × Refugee posts	0.368*** (0.112)	0.339*** (0.093)	0.463*** (0.146)	0.282*** (0.086)	0.864*** (0.264)			
AfD users/Pop. × Refugee comments						1.029*** (0.285)		
AfD users/Pop. × Refugee post likes							0.002*** (0.001)	
AfD users/Pop. × Refugee post share								0.054*** (0.015)
Refugee Attacks/Refugees (t-1)	-0.003 (0.025)							
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) × Posts	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Observations	476,630	480,963	480,963	480,963	480,963	480,963	480,963	480,963
Number of municipalities	4,333	4,333	4,333	4,333	4,333	4,333	4,333	4,333
R <sup>2</sup>	0.046	0.050	0.045	0.052	0.045	0.045	0.045	0.045

*Notes:* This table presents the estimated coefficients from regressions of hate crimes against refugees, scaled by the number of asylum seekers, on social media usage as in Equation (1). *AfD users/Pop.* is the ratio of people with any activity on AfD's Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD's Facebook wall containing the word refugee ("Flüchtling"), divided by 10,000 for readability. Column (1) includes a lagged dependent variable. Column (2) is based on weighted least squares (WLS) based on each municipality's population. In column (3), we replace the number of AfD users calculated over the whole sample with the number of users before the sample start (that is, pre-2015). In column (4), we winsorize attacks at 90th percentile. Columns (5) and (6) present results based on the comments and likes of Facebook posts on the AfD page containing the word refugee, rather than the number of posts on the Facebook wall. Column (7) uses the share of posts containing the word refugee in all posts we observe in a given week. See text for an explanation of the control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Figure A.7: ACCOUNTING FOR THE SKEWED DISTRIBUTION OF AfD USERS**



*Notes:* This figure plots the distribution of the ratio of AfD users in a municipality to population. The red vertical lines indicate the 50th and 90th percentile of the distribution, respectively, which we make use of in Table A.19.

**Table A.19: ACCOUNTING FOR THE SKEWED DISTRIBUTION OF AfD USERS**

Dependent variable AfD user percentiles	Refugee Attacks/Asylum Seekers				
	No zero-user municipalities (1)	Only above median (2)	Only below median (3)	10-90 percentile (4)	User quartiles (5)
AfD users/Pop. $\times$ Refugee posts	0.499*** (0.181)	0.325* (0.187)	0.373** (0.186)	0.711*** (0.259)	
AfD users/Pop. (Q2) $\times$ Refugee posts					0.628 (0.409)
AfD users/Pop. (Q3) $\times$ Refugee posts					2.448*** (0.842)
AfD users/Pop. (Q4) $\times$ Refugee posts					4.971*** (1.158)
Week FE	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes
Observations	395,493	247,863	247,863	345,876	395,493
Number of municipalities	3563	2233	2233	3116	3563
$R^2$	0.046	0.046	0.028	0.040	0.046

*Notes:* This table presents the estimated coefficients from regressions of different hate crimes measures on social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. In column (5), the excluded category is the first quartile of *AfD users/Pop.*; zero-user municipalities are excluded. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.20: ALTERNATIVE TRANSFORMATIONS AND ESTIMATION METHODS**

	Refugee Attack Dummy		Log	Log	Log	
	OLS (1)	Logit (2)	Logit (3)	1+Refugee Attacks (4)	1+Refugee Attacks (5)	1+Refugee Attacks/ Asylum Seekers (6)
AfD users/Pop. $\times$ Refugee posts	0.041*** (0.012)	2.829* (1.517)	2.859* (1.515)	0.029*** (0.009)		
Log(1+AfD users) $\times$ Refugee posts					0.308*** (0.050)	
Log(1+AfD users/Pop.) $\times$ Log(Refugee posts)						0.001*** (0.000)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts	Yes	Yes	Yes	Yes	Yes	Yes
Observations	480,963	130,536	130,536	480,963	480,963	480,963
Number of municipalities	4,333	1,176	1,176	4,333	4,333	4,333
$R^2$	0.081	0.033	0.180	0.118	0.045	0.045

*Notes:* This table presents the estimated coefficients from regressions of different hate crimes measures on social media usage as in Equation (1). The dependent variable is indicated in the top row. *AfD users/Pop.* is the ratio of people with any activity on AfD's Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD's Facebook wall containing the word refugee ("Flüchtling"), divided by 10,000 for readability. D(Refugee Attack) is a dummy equal to 1 if there is an attack in a municipality in a week, and 0 otherwise. In column (1), we estimate this discrete choice model using OLS, in column (2) using conditional fixed-effects logit, and in column (3) using the incidental parameter bias-adjusted two-way fixed effects model of Ferrnandez-Val and Weidner (2016). Log(1+Refugee Attacks) in columns (4) and (5) is the natural logarithm of 1 plus the number of refugee attacks in a municipality in a given week. In column (6), we take the natural logarithm of 1 plus the baseline measure refugee attacks/asylum seekers. See text for an explanation of the control variables. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.21: SPECIFICATION TESTS**

Dependent variable	Residuals Refugee Attacks/Asylum Seekers from Main Regression					
	AfD users/Pop.		Nutella users/Pop.		$I_{Many\ Nutella\ Users}$	
Social media measure	(1)	(2)	(3)	(4)	(5)	(6)
Social media users/Pop. $\times$ Refugee posts (t+1)	0.098 (0.162)	-0.048 (0.233)	0.094 (0.140)	0.279 (0.232)	0.800 (0.771)	1.254 (1.466)
Social media users/Pop. $\times$ Refugee posts (t+2)		0.208 (0.185)		-0.158 (0.181)		0.578 (1.136)
Social media users/Pop. $\times$ Refugee posts (t+3)		0.227 (0.163)		0.058 (0.207)		1.327 (1.071)
Social media users/Pop. $\times$ Refugee posts (t-1)		0.176 (0.259)		0.475 (0.416)		1.295 (2.110)
Social media users/Pop. $\times$ Refugee posts (t-2)		-0.280 (0.197)		-0.310 (0.263)		-2.060 (1.298)
Social media users/Pop. $\times$ Refugee posts (t-3)		-0.204 (0.191)		-0.393 (0.438)		-2.085 (2.071)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	491,260	468,930	491,260	468,930	482,220	468,825
Number of municipalities	4,466	4,466	4,466	4,466	4,466	4,466
$R^2$	0.00	0.00	0.00	0.00	0.00	0.00

*Notes:* This table presents the estimated coefficients from the specification test for the panel regressions. The dependent variable are the residual of the regressions of the ratio of municipality-level refugee attacks to the number of asylum seekers on the contemporaneous interaction term *Refugee posts*  $\times$  *Social media users/Pop.*. For column (1) and (2), we use the residuals from the regression in column (1) of Table 3; for column (3) and (4) the residuals column (3); and for column (4) and (5) the residuals from the regression in column (5) of Table 3. Robust standard errors in all specifications are clustered by municipality. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.22: ROBUSTNESS: MAIN ESTIMATION WITH ALTERNATIVE STANDARD ERRORS**

Dependent Variable Social Media Measure	Refugee Attacks/Asylum Seekers		
	AfD Users/Pop. (1)	Nutella Users/Pop. (2)	$I_{Many\ Nutella\ Users}$ (3)
Social media users/Pop. $\times$ Refugee posts	0.369	0.287	2.317
Clustered by municipality (Baseline)	(0.113)***	(0.091)***	(0.621)***
Clustered by municipality and week	(0.056)***	(0.076)***	(0.472)***
Clustered by county	(0.110)***	(0.088)***	(0.627)***
Clustered by county and week	(0.056)***	(0.076)***	(0.469)***
Clustered by state	(0.131)**	(0.108)**	(0.518)***
Clustered by state and week	(0.054)***	(0.072)***	(0.523)***
Week FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Baseline controls (7) $\times$ Posts	Yes	Yes	Yes
Observations	480,963	480,963	480,852
Number of municipalities	4,333	4,333	4,332
$R^2$	0.045	0.045	0.045

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media usage as in Equation (1). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* and *Nutella users/Pop.* are the ratio of people with any activity on the Facebook pages of the AfD and Nutella, respectively, scaled over the total municipality population. The *Many Nutella Users Dummy* is 1 for municipalities within a county that are in the top tercile of Nutella users per capita, and 0 otherwise. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. Standard errors are calculated as indicated. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.

**Table A.23: COUNTY- AND STATE-LEVEL RESULTS**

	State-Level Results			
	Refugee Attacks/Refugees (1)	Refugee Attacks/Pop. (2)	Log(1+Refugee Attacks) (3)	D(Refugee Attack) (4)
AfD users/Pop. $\times$ Refugee posts	0.891*** (0.092)	0.101*** (0.008)	0.001 (0.002)	0.067*** (0.014)
Week FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Observations	1,776	1,776	1,776	1,776
Number of states	16	16	16	16
$R^2$	0.342	0.457	0.449	0.572

	County-Level Results			
	Refugee Attacks/Refugees (5)	Refugee Attacks/Pop. (6)	Log(1+Refugee Attacks) (7)	D(Refugee Attack) (8)
AfD users/Pop. $\times$ Refugee posts	0.102 (0.117)	0.117*** (0.014)	0.034*** (0.010)	0.045*** (0.011)
Week FE	Yes	Yes	Yes	Yes
Kreis FE	Yes	Yes	Yes	Yes
Observations	44,622	44,622	44,622	44,622
Number of counties	402	402	402	402
$R^2$	0.057	0.088	0.131	0.165

*Notes:* This table presents the estimated coefficients from a regression of hate crimes against refugees on social media usage as in Equation (1), where all variables are collapsed on the state- (“Bundesland”) or county-level (“Kreis”). The dependent variable is the ratio of municipality-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on AfD’s Facebook page to population. *Refugee posts* is the Germany-wide number of posts on the AfD’s Facebook wall containing the word refugee (“Flüchtling”), divided by 10,000 for readability. Robust standard errors in all specifications are clustered by county. \*\*\*, \*\*, and \* indicate statistical significance at the 0.01, 0.05, and 0.1 level, respectively.