

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

Karsten Müller[†] and Carlo Schwarz[‡]

May 24, 2019

Abstract

This paper investigates the link between social media and hate crime. We show that anti-refugee sentiment on Facebook predicts crimes against refugees in otherwise similar municipalities with higher social media usage. To establish causality, we exploit exogenous variation in major Facebook and internet 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 violent crime.

JEL Classification: D74, J15, Z10, D72, O35.

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

*We are grateful to Sascha Becker, Christopher Blattman, Leonardo Bursztyn, Mirko Draca, Thiemo Fetzer, Evan Fradkin, Matthew Gentzkow, Andy Guess, Vardges Levonyan, Atif Mian, Magne Mogstad, Sharun Mukand, Hans-Joachim Voth, Fabian Waldinger, Noam Yuchtman, and seminar participants at the University of Chicago, EEA Conference 2018, Oxford Internet Institute, Geneva Academy of Humanitarian Law, Bruneck Political Economy Workshop, Leverhulme Causality Conference at the University of Warwick, Spring Meeting of Young Economists 2019 and the Royal Economic Society 2019 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 Scholarship from the Leverhulme Trust as part of the *Bridges* program.

[†]Princeton University, The Julis-Rabinowitz Center for Public Policy and Finance, karstenm@princeton.edu

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

1 Introduction

Social media has come under increasing scrutiny in recent years. For example, in the wake of the 2016 presidential election in the United States, relatively recent phenomena such as fake news, social media echo chambers, and bot farms have been subjects of widespread media coverage and public discourse (e.g New York Times, 2016, 2017a). The role of online hate speech in particular has been at the center of an intense and polarized debate. Despite public interest and calls for policy action, there is little empirical evidence on how social media hate speech translates into real-life behavior.

Building on the literature on media exposure and violence, we investigate the role of social media in the propagation of hate crime. Previous research has shown that traditional media can play a role in violent outbursts or ethnic hatred (e.g Yanagizawa-Drott, 2014; Adena et al., 2015; DellaVigna et al., 2014). In contrast to traditional media, social media platforms allow users to self-select into preferred topics and viewpoints. This preferential selection may limit the spectrum of information people absorb and create “echo chambers” (Sunstein, 2009, 2017), which reinforce similar ideas (see e.g. Bessi et al., 2015; Del Vicario et al., 2016; Schmidt et al., 2017). Social media has also become a widely-consumed news source, particularly for young people: in Germany, for example, social media is among the main news sources of 18 to 25 year olds (Hölig and Hasebrink, 2016). In the US, around half of all adults use social media to get news and two thirds of Facebook users use it as a news source (Pew Research Center, 2018). This suggests that social media could be particularly effective in propagating hateful sentiments.

In this paper, we study the link between social media and hate crime by 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 recent influx of around one million refugees into the country between 2015 and 2016 (BAMF, 2016) and the unfortunate frequency of violent crimes committed against them (see, for example, recent video coverage by New York Times, 2017b). Between January 2015 and early 2017 alone, the non-profit organization “Amadeu Antonio Stiftung” recorded around 3,300 anti-refugee incidents, including almost 750 cases of arson or outright assault.

We posit that social media information channels can reinforce anti-refugee sentiments, which may push some potential perpetrators over the edge to carry out violent acts. If social media plays a role, we would expect more hate crimes to occur in municipalities with higher

exposure to Facebook, particularly when tensions are high.

Our empirical strategy thus exploits differences in Facebook usage at the municipal level and weekly variation in anti-refugee salience on social media. We create a novel 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 420,000 followers, their Facebook page also has a broader reach than that of any other German party (see Appendix A for a history of the AfD).

This widespread reach makes the AfD’s Facebook page uniquely suited to measure anti-refugee sentiment on social media. 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. The AfD 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 generated by over 93,000 individual users, our data provide a unique insight into far-right Facebook usage. Based on this detailed data, we further construct a measure of a municipality’s exposure to Germany-wide anti-refugee sentiment using the share of the population that is active on the AfD Facebook page.

By combining these two proxies for social media activity, we find that anti-refugee hate crimes disproportionately increase in areas with higher Facebook usage during periods of high anti-refugee salience. This correlation is especially pronounced for violent incidents against refugees, such as arson and assault. Controlling for a large vector of municipality characteristics, interacted with our salience measure, makes little difference for the magnitude and statistical significance of these estimates.

The concern, however, is that our measures of Facebook usage could correlate with unobservable municipal characteristics that explain the disproportionate increases in hate crimes during times of high anti-refugee sentiment. To narrow down the social media transmission channel, we provide quasi-experimental evidence using country-wide Facebook outages and local internet disruptions. Both of these induce plausibly exogenous variation in exposure to social media and thus should allow us to retrieve causal estimates.

To begin, we study large, Germany-wide Facebook outages resulting from programming

or server problems at the platform. These outages disrupt users’ exposure to this particular social media platform without affecting other online channels. Consistent with a causal effect of social media, we find that Facebook disruptions reduce local hate crimes, particularly in areas with many AfD users. Further, during Facebook outages, higher anti-refugee sentiment is not associated with a differential increase in hate crimes in areas with high Facebook usage. These results suggest that social media can play a propagating role in translating online content into offline violence.

We also exploit hundreds of local internet disruptions as a source of highly granular exogenous variation in access to social media. In particular, we identify municipalities and weeks with severe internet outages, which we verify using news articles in regional and national media outlets. These *local* internet disruptions reduce municipalities’ exposure to hate speech on social media while leaving *Germany-wide* anti-refugee salience unaffected. Notably, internet disruptions are geographically dispersed and orthogonal to both AfD likes on Facebook and the total number of refugee attacks in a municipality.

We find that, while the number of hate crimes increases during periods of higher anti-refugee salience, this correlation disappears for municipalities experiencing an internet outage. Quantitatively, a typical internet disruption *fully* mediates social media’s link with hate crime. Further, internet outages themselves do not appear to affect the number of anti-refugee incidents beyond their impact through social media usage. This makes it unlikely that we are capturing a “displacement effect” which could arise if potential perpetrators were merely busy fixing their internet access. This further points to social media as the propagation mechanism. Additionally, we do not appear to be capturing other online channels: internet outages have *no* mediating effect on hate crimes in areas with higher *general* internet usage, once we take social media usage into account.

To further investigate social media’s capacity as a transmission channel, we show our results also hold for alternative measures of Facebook usage. More specifically, we create proxies for municipal-level Facebook activity based on the number of users on the “Nutella Germany” page.¹ With over 32 million likes, Nutella has one of the most popular Facebook pages in Germany and therefore provides a measure of *general* social media usage, independent of support

¹Nutella is a popular chocolate-hazelnut spread that is consumed by people from all socioeconomic backgrounds and all regions. This broad popularity makes Nutella’s Facebook page more advantageous for our analysis than other popular pages such as those for FC Bayern Munich or BMW, whose users may be clustered around a particular geographical area or socioeconomic group.

for the AfD.² We show that municipalities with many Nutella users per capita also experience more anti-refugee incidents in times of high refugee salience on the AfD page – unless social media access is disrupted by Facebook or internet outages.

The link we uncover appears to be specific to anti-refugee sentiment: other posts on the AfD Facebook page, e.g. those related to Muslims or the European Union, do not have significant predictive power for anti-refugee hate crimes. Consistent with the hypothesis that social networks can act as transmission channel, the correlation with hate crime is larger in regions where AfD users show higher Facebook engagement via likes and comments. Importantly, these engagement proxies are uncorrelated with social media usage and thus provide meaningful additional variation. We also analyze how other salient news events mediate the link of anti-refugee Facebook posts with the number of violent incidents, building on Eisensee and Strömberg (2007) and Durante and Zhuravskaya (2018). 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. During weeks in which these events partially crowd out the salience of refugees, social media has a significantly more muted relationship with hate crimes.

When interpreting our results, we do not claim that social media itself causes crimes against refugees out of thin air. Rather, our argument is that social media can act as a propagating mechanism for hateful sentiments that likely have many fundamental sources: local differences in xenophobic ideology or immigrant prevalence are only two obvious examples. The evidence in this paper suggests that quasi-random shifts in exposure to anti-refugee sentiment on social media can increase the number of anti-refugee attacks.

Related literature. Our work provides evidence that social media may have propagating effects on real-life outcomes, as measured by hate crimes. We build on existing work on the impact of media exposure and persuasion (see e.g. DellaVigna and Gentzkow, 2010; DellaVigna and La Ferrara, 2015). In addition to the already cited work on traditional media and violence, 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 due to displacement effects. Television has also been associated with short-lived outbursts of domestic violence (Card and Dahl, 2011). In other research, Bhuller et al. (2013) demonstrate that exposure to pornographic material on the internet is linked to increased sex crime. Bursztyn

²We also use the Nutella data to create a dummy for municipalities with many Facebook users *within a county*. This measure only correlates with municipality characteristics that reflect patterns of social media usage in Germany (also see Section 3.5.1).

et al. (2017) find that media coverage of close elections increases voter turnout, while Gavazza et al. (2018) show that broadband diffusion decreased voter turnout in the United Kingdom (see also Gentzkow, 2006; Manacorda et al., 2017). Furthermore, Enikolopov et al. (2016) argue that social media exposure spurs protest participation in Russia by reducing coordination costs.

We contribute to this literature by investigating the social media’s role in violence. Previous research has already documented the high prevalence of exposure to hate speech online (Oksanen et al., 2014). Other work has shown that Google search data can be used to measure racial animus (Stephens-Davidowitz, 2014). Müller and Schwarz (2018) document correlations between Twitter usage and hate crimes in the US. We show that hateful sentiments are not only propagated through social networks but also increase the occurrence of hate crimes.

Our paper also builds on research about the polarization of citizens (e.g Fiorina and Abrams, 2008). There is no consensus in the existing literature on whether social media increases or decreases polarization: some authors argue that social media are divisive (Pariser, 2011; Gabler, 2016), while others find that polarization *decreases* with social media usage (Boxell et al., 2017; Barberá, 2014). Our work suggests that even if overall polarization is unaffected by social media, the content in online networks can be associated with violent crimes.

Additionally, we contribute to the literature on culture and violence. Summarizing a vast body of research, Alesina and La Ferrara (2005) find that cultural and religious fragmentation 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 of sentiment leading to violent incidents. As such, our work is also 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 align 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, legislators are now beginning to address social media. Our work is thus particularly topical in light of the political discussions in many countries about anti-hate speech laws and censoring hate speech on social media. The German parliament, for example, passed an anti online hate speech

law (“Netzwerkdurchsetzungsgesetz”) on June 30, 2017, which threatens providers of online platforms such as Facebook with fines up to EUR 50 million for failing to delete “criminal” content that is “obviously unlawful”. The controversial law was the initiative of German Minister of Justice Heiko Maas, who lamented social media platforms’ unwillingness to address “online hate crime”.³ The European Union has issued independent guidelines calling on social media companies to remove illegal hate 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 presents basic correlations in the data, the empirical strategy, and the main results. Section 4 concludes.

2 Data

We construct a new data set comprising 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 12 different sources which we describe in more detail in the following subsections: (1) Municipal-level data on anti-refugee hate crimes; (2) Facebook data on posts, likes, and comments on the AfD pages; (3) hand-collected municipal-level data on Facebook user locations; (4) municipal-level data on internet outages; (5) a hand-coded dataset on major weekly Facebook outages; (6) municipal- and county-level socioeconomic data from the German Statistical Office; (7) election district voting data; (8) county-level data on broadband access; (9) Eurostat survey data on internet usage; (10) municipal-level data on newspaper sales; (11) city-level data on neo-Nazi murders and historical anti-Semitism; and (12) 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 and Table 2. The online appendix provides a comprehensive overview of the data sources and variable definitions (see Table A.5).

³See, for example, the official statement of the German parliament on bundestag.de.

Table 1: Summary Statistics for Main Variables

	Level	Obs.	Mean	S.D.	Min.	Max.
Refugee Attacks[†]						
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
Social Media Data						
AfD users/Pop. [†]	Municipality	495,726	3.006	2.863	0	80.25223
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
Auxiliary Variables						
$I_{Internet\ outage}$	Muni.-Week	495,726	0.001	0.025	0	1
$I_{Facebook\ outage}$	Week	495,726	0.081	0.273	0	1
Nutella users/Pop. [†]	Municipality	495,726	2.007	2.040	0	35.82689
$I_{Many\ Nutella\ Users}$	Municipality	495,726	0.417	0.493	0	1
Baseline Controls						
Population (2015) [†]	Municipality	495,726	1.840	7.478	0.034	352.003
GDP/Worker	County	493,617	63,095	9,846	46,835	136,763
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
Raw Data						
Refugee attacks	Muni.-Week	495,726	0.007	0.099	0	8
Refugees (2015) [†]	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 main variables in the estimation sample. Variables tagged with a [†] are scaled by population in 10,000. Share variables are in percent.

2.1 Anti-Refugee Incidents

The data on incidents targeting refugees were collected by the Amadeu Antonio Foundation and Pro Asyl (a pro asylum NGO).⁴ These data cover anti-refugee activity including anti-refugee graffiti, arson of refugee homes, assault, and incidents during protests in Germany between January 2015 and early 2017. This period is of particular interest since it includes the beginning and height of the refugee crisis in Germany. All 3,335 anti-refugee aggressions feature a short description and are classified into four groups. The most common occurrences are property damage to refugee homes (2,226 incidents), followed by assault (534), incidents during anti-refugee demonstrations (339), arson (225). 11 events are classified as suspected cases that were still under investigation. Table A.1 in the online appendix list examples for each class of anti-refugee activity.

All incidents are geo-coded with an exact longitude and latitude, which we use to assign them to municipalities.⁵ Figure 1 shows the total number of anti-refugee events per number of asylum seekers in our observation period for each German municipality.

The data appear to be high quality, with each entry including its source. Nearly half of the incidents in the dataset are reported by the federal government in response to inquiries by the left-wing party “Die Linke”. Other sources include police reports and national or local media outlets. We hand-checked a random sample of 100 incidents and found their coding accurately reflected the information reported in the respective source. Note that our measures of anti-refugee activity are necessarily right-skewed, because we use a municipality-week panel; we discuss a host of strategies to address this in Section 3.6.

2.2 Facebook Data on Refugee Salience

We use the AfD’s Facebook page to construct a measure for the salience of anti-refugee hate speech on social media. We chose the AfD’s page because the AfD is by far the largest far-right party in Germany and has the highest number of Facebook followers of *any* German party. This

⁴These data are available at <https://www.mut-gegen-rechte-gewalt.de/service/chronik-vorfaelle>.

⁵To assign coordinates to municipalities, we use the shape files provided by the ©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 any 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 than the 11,165 German “Gemeinden” and are therefore more suitable for spatial analysis according to the data provider (see link).

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.

Figure 2: AfD Facebook Usage per Capita, by Municipality



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 (AfD).

makes the page arguably the most important platform of exchange about refugees for Germany’s right-wing social media users.

We start by using the Facebook Graph API to collect all status posts, comments, and likes from the AfD Facebook page (see Appendix B.1. for an introduction to Facebook). The API provides a unique identifier for each post, allowing us to link posts to comments and likes, as well as the users that posted, commented, or liked 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.

As our baseline measure for anti-refugee hate speech salience on social media, we use the number of posts on the AfD Facebook page that contain the word “Flüchtling” (refugee) in any given week. We also construct analogous measures based on comments or likes.

A potential downside of this procedure is that we may inadvertently tag posts that do not express negative sentiments towards refugees. However, a careful content analysis of posts and comments reveals that the overwhelming majority appear to agree with the positions of the AfD. This is perhaps unsurprising given that only people who “like” the AfD Facebook page will be informed about new posts: critics, on the other hand, have a strong incentive not to indicate publicly that they “like” the party.

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?”⁶ Table A.2 in the online appendix, includes further representative examples of posts published on the AfD page. We also construct measures for the salience of other topics by tagging AfD Facebook page posts containing the words “Islam”, “Muslim”, “Jude” (Jew) or “EU”.

2.3 Municipal-Level Facebook Usage Data

For our empirical strategy, we construct a measure of social media usage at the municipal level. Because survey data on German 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.

⁶Original Post in German: “Evtl. soll ja die deutsche Genetik ausgerottet werden, durch große Flüchtlingsströme. Doch was könnte der Grund sein, WW2?”

As we are interested in the transmission of right-wing social media sentiment, we measure Facebook usage based on users of the AfD Facebook page. In total, we can identify 93,806 users who interacted with the page at least once.⁷ We were able to hand-collect a place of residence for 39,632 of these users. Overall, we were able to identify at least one AfD Facebook page user for 3,565 of the 4,466 municipalities. Note that the decision of users to disclose their location is unlikely to matter in our setting. This is because we exploit variation *within* the same location over time, which abstracts from time-invariant endogenous selection using municipality fixed effects.

Using the location data for AfD users, we can also assign posts, comments, and 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 therefore 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.

For robustness, we also create an alternative measure of *general* Facebook usage in Germany based on the page of Nutella Germany. We use the Nutella page because, with more than 32 million followers, it is one of the most popular Facebook page in Germany (Focus, 2014) and should therefore provide a reasonable proxy of Facebook activity across municipalities. 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. Compared to the AfD’s Facebook page, we have considerably less user data for Nutella’s Facebook page despite its much higher number of fans because we can only collect data for users who 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 general Facebook usage in Germany and a considerable improvement over existing survey measures. We have at least one Nutella user for 3,190 municipalities. Based on the Nutella data, we can also create a dummy variable equal to 1 for municipalities in the top tercile of Nutella users per capita *within a county* ($I_{Many\ Nutella\ Users}$).

⁷The Facebook API does not provide data on which users “like” a page but only on users who *interact* with a page, e.g. by liking another user’s comment. As a result, the total number of user IDs we have is smaller than the more than 300,000 people who had liked the AfD Facebook page at the time of data collection.

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 telephone area codes and includes start times and durations. We use the area codes to assign internet problems to municipalities; the start date and duration allow us to count each municipality’s number of problems in each week.⁸ The internet outage reports are geographically dispersed with no clear patterns of regional clustering (see Figure A.3).

To validate the Heise data, we search for newspaper reports on major internet disruptions. For all major disruptions we could identify in newspapers, the data suggest an increase in the number of outages specific to the internet provider experiencing the outage. Table A.3 lists several examples of newspaper reports on such outages and the respective information in our data.⁹

Because some reports may reflect individual users’ glitches rather than general disruptions, we exclude reported outages with a duration of less than 24 hours.¹⁰ 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. Note that we scale outages over population because towns with more inhabitants mechanically also report more disruptions. As we will discuss later, our results are robust to using alternative definitions of this cut-off.

In addition to 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 eight large outages (see Table A.4 for an overview and more details). For independent validation, we also obtain the number of weekly user-reported Facebook problems on the Facebook page of “Allestörungen”, a portal for aggregating user complaints on individual websites and apps. Perhaps unsurprisingly, the eight outages widely reported on in the news media are also associated with spikes in user-reported problems.

⁸If an area code spans multiple municipalities, we assign an internet outage to the municipality that overlaps most with the area code. This is preferable to assigning the outage to all municipalities within the area code’s territory because some area codes include minor overlaps with many municipalities. Assigning the internet outage to all of these municipalities would introduce substantial noise.

⁹To 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.

¹⁰In 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.

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 in a given week from our Heise data. The downside is that Facebook outages 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. We include information on each municipality’s population 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 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, 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, information collected by the Federal Ministry of Transport and Digital Infrastructure. Broadband access is highly correlated with publicly available survey data, from Eurostat, on individuals’ internet use; these data are only available on the state level (see Figure A.4 in the online appendix).¹¹ 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 2016/2017 newspaper sales 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 (1) sell at least 50 copies and (2) have a market share of at least

¹¹More specifically, 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.

1%. Based on this data, we construct a measure of traditional newspaper consumption as the number of newspaper sales per capita.¹²

To measure 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.¹³

Finally, we obtain Google trends data on 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 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).

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 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 AfD Facebook page posts about refugees against the number of anti-refugee incidents in Figure 3a. Weeks with more anti-refugee posts also tend to

¹²To have approximately equivalent samples across specifications, we impute values for 1,013 towns for which news paper sales data are not available, based on a municipality’s population, population density, AfD vote share, and county fixed effects. However, the results are almost equivalent without imputation, which leaves us with a somewhat smaller sample of 3,320 municipalities (available upon request).

¹³Note that we use the natural logarithm of one plus the 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 Stürmer letters as zero.

have more anti-refugee events. The correlation also holds in a time series regression of refugee attacks on AfD posts, which yields a R^2 of around 34% (reported in online appendix Table A.7). 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 one refugee attack, depending on whether we can identify *at least one* AfD Facebook page user. Municipalities with AfD users are three times as likely to experience an attack during our observation period. Out of the total 3,335 attacks on refugees in our sample, 3,171 occurred in municipalities with AfD Facebook page 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-run the time series regressions separately for municipalities with at least 1 user on the AfD Facebook page and contrast it with municipalities with 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 correlation with 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.7 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 suggest a strong statistical 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 correlation is driven by users on Facebook *reacting* to, rather than causing violence. It is also unclear to what extent common shocks such as the news cycle or persistent local attitudes towards minorities affect both the online and offline behavior we observe. 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 link between social media and anti-refugee incidents, we estimate fixed effects panel regressions akin to a Bartik-type approach (Bartik, 1991). In particular, we use the interaction of local Facebook usage ($Social\ Media\ Users/Pop_m$) and refugee posts on the AfD Facebook page ($Refugee\ Posts_w$) to measure the differential change of hate crimes conditional

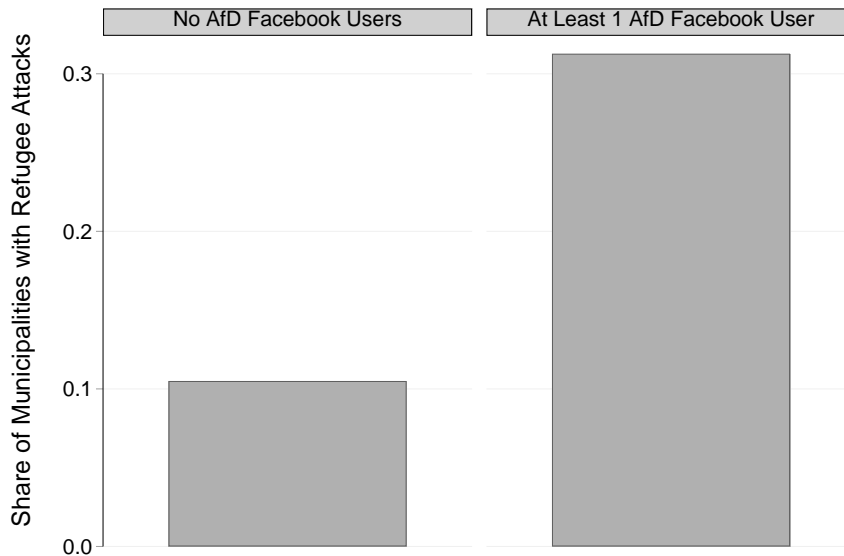
Figure 3: Introductory Correlations

(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 Facebook page user in the municipality. We are able to identify at least one AfD user in 3,563 municipalities; for 903 municipalities we find no AfD user.

on anti-refugee sentiment on social media. This empirical set-up creates variation by week and municipality, which we exploit in the following regression model:

$$\begin{aligned}
Attacks/Refugees_{mw} = & \beta \text{ Social Media Users/Pop}_m \times \text{Refugee Posts}_w \\
& + \gamma \text{ Controls}_m \times \text{Refugee Posts}_w \\
& + \text{Week FE}_w + \text{Municipality FE}_m + \epsilon_{mw},
\end{aligned} \tag{1}$$

In this specification, β measures the differential change in anti-refugee incidents conditional on Germany-wide posts about refugees on the AfD page – as a proxy of Germany-wide anti-refugee salience on social media – and social media users per capita.¹⁴ Our baseline measure of local exposure to right-wing social media is *AfD Users/Pop.*, the share of people active on the AfD Facebook page. Because attacks are a function of the number of refugees in a given area, we scale them over the number of asylum seekers. In the robustness section, we explore a plethora of different variable transformations and show our results are highly robust. We cluster standard errors by municipality; in Section 3.6, we show that clustering on other levels, if anything, makes our results more statistically significant.¹⁵

Our framework has three advantages. First, one out of the 4,466 municipalities cannot induce meaningful variation in the Germany-wide weekly *Refugee Posts_w* measure.¹⁶ As a 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.¹⁷ Third, the panel format allows us to abstract from unobserved weekly factors and municipal-level predictors of attacks using a full set of fixed effects.¹⁸ Week fixed effects absorb changes in the number of anti-refugee incidents that affect all municipalities to the same extent, e.g. nationwide news events on increased

¹⁴We also experimented with many alternative ways to measure anti-refugee salience and social media usage, all of which yield similar results. See table A.14 and table A.16 in the online appendix.

¹⁵In the online appendix, we show that the findings also hold on the county- and state-level.

¹⁶Note 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 the weighted sum of the individual geographical areas' growth rates and hence one uses a leave-one-out measure of industry growth rates. Due to the differences 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.

¹⁷In 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 for whom we have location data. It turns out that this adjustment makes little difference for the results.

¹⁸Note that the non-interacted terms for refugee posts and local users are absorbed by the fixed effects.

refugee numbers. 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.

A concern with estimating equation 1 is that $AfD\ Users/Pop.$ may be correlated with other municipality characteristics that explain the salience measure’s link with anti-refugee attacks. For example, the share of AfD Facebook subscribers could be a mere proxy for a local right-wing presence, which could lead to more anti-refugee attacks in times of high anti-refugee salience. While including interactions of a plethora of municipality characteristics listed in Table 1 and Table 2 with our anti-refugee salience measure does not substantially affect the size or statistical significance of the estimated coefficient of β , this alone does not allow us to identify causal effects.

We therefore develop an identification strategy based on Facebook and internet outages. These disruptions induce quasi-experimental variation in people’s exposure to social media while leaving Germany-wide anti-refugee salience and local characteristics unchanged. Goldsmith-Pinkham et al. (2018) argue that the exogeneity of shock exposure – in our setting, social media usage – is a sufficient condition for retrieving causal effects in a Bartik framework. Therefore, the quasi-random reduction in social media exposure from outages should allow us to derive causal estimates.

The first part of this empirical strategy exploits the timing of major server problems at Facebook, which disrupt access to the platform. These disruptions are a recurring phenomenon (see Table A.4 for a list of newspaper references). In the second part, we build on the insight that the German internet infrastructure was trailing behind that of many other European Countries (e.g. Latvia) and the OECD average (see Financial Times, 2017; OECD, 2016). As a result, prolonged internet outages are more common. Because around 50% of worldwide Facebook users accessed the platform with their computers, many users are exposed to disruptions in internet access. In Germany, this share is likely to be even higher because of the relatively slow adoption of mobile internet.¹⁹

As described in Section 2, the data on internet outages include disruptions resulting not only from physical infrastructure damages but also provider-specific problems; in fact, many of the large outages originated from problems at a single local provider. These outages are widely geographically dispersed: Figure A.3 visualizes the distribution of disruptions per capita across Germany. Crucially, the frequency of internet problems is uncorrelated with our social media

¹⁹Data on Facebook usage patterns reported on Statista.com and on mobile internet usage in Germany on (also on Statista.com) support this assessment.

measures; as such, internet disruptions provide exogenous variation that is not already captured by our variables on local Facebook usage. The number of reported internet problems is also uncorrelated to the total number of refugee attacks in a given municipality ($t - stat = -1.34$).

If Facebook and internet outages indeed reduce municipalities' exposure to social media content, we would expect them to weaken the relationship between refugee posts and anti-refugee attacks. Put differently, during outages, we should observe fewer anti-refugee attacks at times of high salience in municipalities with high Facebook penetration. These outages also allow us to further address the concern of reverse causality. If we were merely capturing that incidents drive posts on social media, Facebook and internet outages should not reduce the number of hate crimes. Instead, they should only reduce social media activity, keeping the number of anti-refugee incidents unchanged.

The underlying identifying assumption of this approach is that Facebook and internet outages only affect anti-refugee incidents through their effect on social media exposure. This assumption appears highly plausible for Facebook outages. In the case of internet outages, for which we have highly granular variation at the municipality-week level, one may be worried about alternative online channels. We discuss potential threats to identification in more detail in the next section.

We analyze the effect of Facebook and internet outages in a triple difference framework. The former only vary by week, the latter by municipality and week. For the internet outages, we estimate the following regression model:²⁰

$$\begin{aligned}
Attacks/Refugees_{mw} = & \beta \text{ Social Media Users/Pop}_m \times \text{Refugee Posts}_w \\
& + \lambda \text{ Outage}_{mw} \times \text{Social Media Users/Pop}_m \times \text{Refugee Posts}_w \\
& + \delta_1 \text{ Outage}_{mw} + \delta_2 \text{ Outage}_{mw} \times \text{Refugee Posts}_w \\
& + \delta_3 \text{ Outage}_{mw} \times \text{Social Media Users/Pop}_m \\
& + \gamma_1 \text{ Controls}_m \times \text{Refugee Posts}_w \\
& + \gamma_2 \text{ Controls}_m \times \text{Outage}_{mw} \\
& + \text{Week FE}_w + \text{Municipality FE}_m + \epsilon_{mw},
\end{aligned} \tag{2}$$

²⁰Note that, since internet outages neither induce changes in the Germany-wide anti-refugee salience (Refugee Posts_w) nor have an effect on our time-invariant $\text{Social Media Users/Pop}_m$ measures, it is not possible to use internet outages as an instrument in our setting. However, our triple-difference set-up does create exogenous variation in exposure, and is therefore related to an instrumental variable estimation.

Table 2: Summary Statistics for Additional Controls

	Level	Obs.	Mean	S.D.	Min.	Max.
Additional Media and Internet Controls[†]						
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	491,175	924.857	768.374	0	16442.730
Additional Right-Wing Controls						
Nazi murders/Pop. [†]	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
Additional Socio-Economic Controls						
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
Additional Voting Controls (2017 Election)						
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
Additional Demographic Controls						
Share aged 0-24	Municipality	495,726	0.239	0.054	0	0.371
Share aged 25-49	Municipality	495,726	0.323	0.062	0	0.454
Share aged 50-74	Municipality	495,726	31.533	6.513	0	50.080
Share aged 75 and older	Municipality	495,726	9.037	2.424	0	17.649

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.

The main coefficient of interest λ captures the correlation of anti-refugee attacks and Facebook usage, depending on whether an outage occurs in a given municipality in a given week. To facilitate the interpretation of the individual coefficients, we will include the interactions one by one. In particular, we start by only estimating the effect of the outage dummy. Then, we provide evidence that its effect is driven by areas with high social media usage at times of high anti-refugee sentiment. The vector $Controls_m \times Outage_{mw}$ controls for these outages' differential effects on municipalities based on their observable characteristics, of which general internet affinity is the most important.

Because Facebook outages only vary by week, we replace $Outage_{mw}$ with $Outage_w$ (in Equation (2)); as a result, the estimates of δ_1 and δ_2 in Equation (2) are absorbed by the week fixed effects. The triple interaction $Outage_w \times Refugee Posts_w \times Social Media Users/Pop_m$ measures if the differential effect of anti-refugee salience between high and low Facebook usage municipalities disappears in weeks with Facebook outages.


3.3 Baseline Results

We first illustrate the intuition behind our regression framework in Figure 4. The figure shows a binned scatter plot of anti-refugee attacks conditional on anti-refugee salience, split by the degree of Facebook usage. While higher anti-refugee salience is associated with more anti-refugee attacks in both samples, the increase is far more pronounced for municipalities with high Facebook usage (Panel (a)). Our baseline regression coefficient picks up the difference in slopes between municipalities based on local Facebook usage.

Table 3 presents the regression results from estimating Equation (1) with varying sets of control variables (interacted with refugee salience). The coefficient of the interaction of local Facebook usage and Germany-wide refugee posts is positive and highly significant in all specifications. Column (1) shows the panel regressions with the baseline control variables. For brevity, we report the full set of estimated coefficients for the baseline controls in the online appendix (Table A.9).

Taken at face value, the coefficient of 0.369 on the interaction term indicates a large magnitude. Consider weeks with particularly high sentiment (one standard deviation above the mean) compared to those with low sentiment (one standard deviation below the mean). This difference is associated with about 123 refugee posts ($2 \times 61.42 \approx 122.844$). As a case study, consider the cities of Bochum and Hannover, which are about one standard deviation apart in AfD users over population (≈ 2.9). The estimated magnitude of local exposure to heightened

Figure 4: Exposure to Anti-Refugee Sentiment and Hate Crimes



Notes: This figure plots the average number of anti-refugee attacks against our measure of anti-refugee sentiment for municipalities below and above the median of *AfD Users/Pop*. Refugee attacks are binned by 20 quantiles of refugee posts. The number of anti-refugee attacks was residualized with respect to municipal population.

anti-refugee sentiment is given by $0.369 \times 122.844 \times 2.86$ (in 10,000) ≈ 0.013 .²¹ This correlation implies 35% more hate crimes relative to the average number of 0.038 refugee attacks per 10,000 asylum seekers.²²

Importantly, this correlation does not appear to be driven by support for the AfD: the interaction of anti-refugee posts with the AfD vote share in the 2017 federal election is not statistically significant. In other words, after accounting for the presence of AfD Facebook users, higher anti-refugee salience is not associated with increases in refugee attacks in municipalities with more AfD voters. This highlights a distinction between our social media measure and general support for the party.

We next address observable differences between municipalities with higher shares of AfD Facebook users by introducing a richer set of controls in the following columns (see Table 2 for an overview of the control variables). In column (2), we first include additional right-wing controls, namely the vote share of the neo-Nazi party NPD and proxies for anti-minority violence. The interacted controls hardly have 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, in column (3), we investigate if our result is driven by general media exposure. We find 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) similarly has little effect.²³ In column (6) we include further flexible controls for the local age structure. Again, conditioning on additional controls makes no discernible difference for the point estimates.

The only noticeable effect on our results comes from the joint inclusion of all interacted controls in column (7). This reduction is likely explained by the fact that part of the across-group variation in social media usage is absorbed by the use of 29 control variables interacted with refugee posts. Taken together, these results are a first indication that the correlation between social media exposure and anti-refugee incidents is not driven by observable municipality differences unrelated to Facebook usage.

We also explore heterogeneity across different types of anti-refugee incidents in table

²¹Note that the coefficients in Table 3 are multiplied with 10,000 for readability.

²²Note that the small average 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.

²³In unreported regressions, we find that the results are nearly identical if we control for the election results of the 2013 federal election instead.

Table 3: Baseline Correlations: Facebook Posts and Hate Crime

	Additional interacted controls						
	Baseline controls (1)	Right-wing controls (2)	Media controls (3)	Socio-economic controls (4)	2017 vote controls (5)	Age structure controls (6)	All controls (7)
AfD users/Pop. \times Refugee posts	0.369*** (0.113)	0.360*** (0.111)	0.325*** (0.105)	0.355*** (0.112)	0.346*** (0.108)	0.346*** (0.113)	0.285*** (0.098)
Observations	480,963	480,963	480,963	475,302	480,963	480,963	475,302
Number of municipalities	4,333	4,333	4,333	4,282	4,333	4,333	4,282
R^2	0.045	0.045	0.045	0.046	0.045	0.045	0.046
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) \times Posts	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Right-wing controls (4) \times Posts		Yes					Yes
Media controls (3) \times Posts			Yes				Yes
Socio-econ. controls (4) \times Posts				Yes			Yes
Election controls (7) \times Posts					Yes		Yes
Age controls (4) \times Posts						Yes	Yes

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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. All control variables are interacted with the *Refugee posts* measure; see text for a description of the controls. Robust standard errors in all specifications are clustered by municipality. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

Table 4. We find that the correlations are entirely driven by more violent crimes, namely arsons, assaults, and miscellaneous property damage. We do not find a significant link for incidents during 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. This result also adds to the findings in Enikolopov et al. (2016), who find that social media spurs protest participation in Russia.

Table 4: Social Media: Violent vs. Non-Violent Incidents

	Type of Incident				
	Arson (1)	Property damage (2)	Assault (3)	Demonstrations (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 the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers; see A.5 for definition of attack types. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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. All control variables are interacted with the *Refugee posts* measure; see text for a description of the controls. Robust standard errors in all specifications are clustered by municipality. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

Overall, the findings we present in this section are suggestive of a strong correlation between exposure to anti-refugee salience on social media and the number of violent attacks on refugees. Next, we exploit quasi-random variation in Facebook and internet outages in an attempt to establish causal effects.

3.4 Quasi-Experimental Evidence: Facebook and Internet Outages

To isolate the importance of social media, we next draw on internet and Facebook outages as sources of quasi-experimental variation. Our baseline measure of internet outages is a dummy for values in the top quartile of internet disruptions per capita, which vary by municipality and

week (see section Section 2 for more details). This gives us 313 severe internet outages.²⁴ The eight major Facebook outages we exploit only vary by week. We discuss the results and their interpretation in turn.

Internet outages. Are local internet outages severe enough to decrease a municipality’s exposure to social media? We investigate this question using our data from the AfD Facebook page. To proxy for local social media activity, we use a sample of posts and comments for which we know the users’ locations.²⁵ Figure 5 plots the mean of these variables and the corresponding 95% confidence intervals, divided by whether a municipality experiences an outage in a given week. In both cases, we observe significantly less local activity on Facebook for municipalities that experience an internet disruption. In Table A.10 in the online appendix, we report a range of regression-based tests that yield similar results. This lends credence to the idea that exposure to social media content is reduced in the affected municipalities and not made up for by users accessing Facebook with their mobile phones.

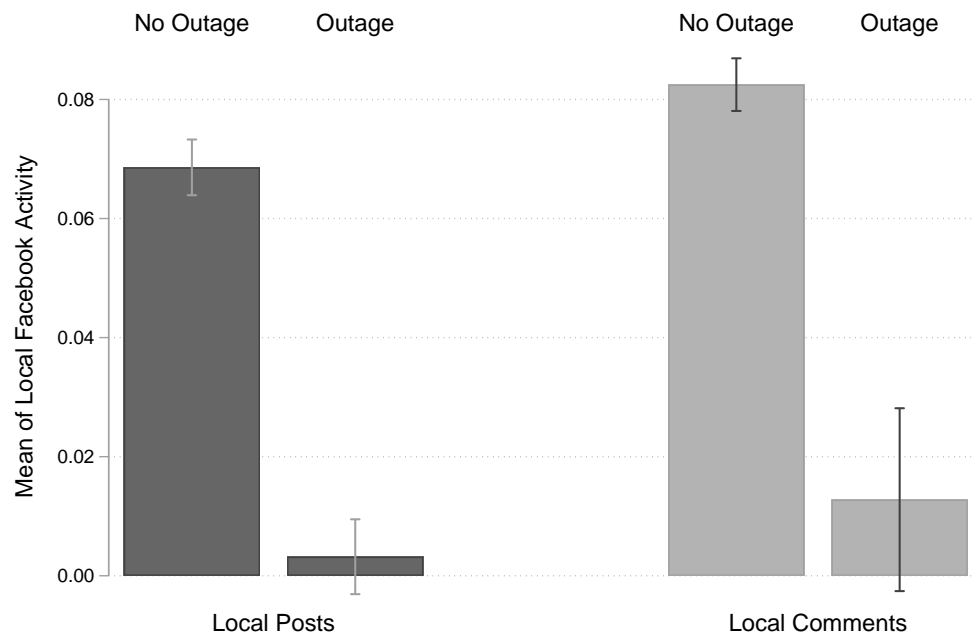
If internet outages indeed reduce local social media exposure, we would expect them to mediate social media’s capacity to propagate anti-refugee incidents. As described in Section 3.2, we test this hypothesis by interacting the main terms of interest $Social\ Media\ Users/Pop_m \times Refugee\ Posts_w$ with $Internet\ Problems_{mw}$, our dummy for severe internet disruptions. We graphically illustrate the results in Figure 6. The binned scatter plot is identical to Figure 4, except that we separately plot attacks for municipalities that experience an internet outage. This reveals a striking pattern: while anti-refugee attacks increase with anti-refugee posts, this relationship disappears completely in municipalities that experience an internet outage. This holds true for municipalities with both high and low Facebook usage.

Figure 6 implies that the attenuation effect of internet outages is substantial. Consider the correlation of refugee posts and attacks in panel (a). Without outages, the individual data points of this correlation almost lie on a 45 degree line. During outages, the correlation is essentially zero. This means that the outage effect is larger than the baseline estimate of $AfD\ Users/Pop. \times Refugee\ posts$, which is given by the slope difference of the dotted lines in panels (a) and (b). Our interpretation is that this is because outages completely cut off users from social media.

²⁴In the online appendix, we show our results are robust to both alternative definitions and using the continuous ratio of the local disruptions per capita (which also varies by week). In the latter case there are more than 1200 municipality-week observations with internet problems (see Table A.12).

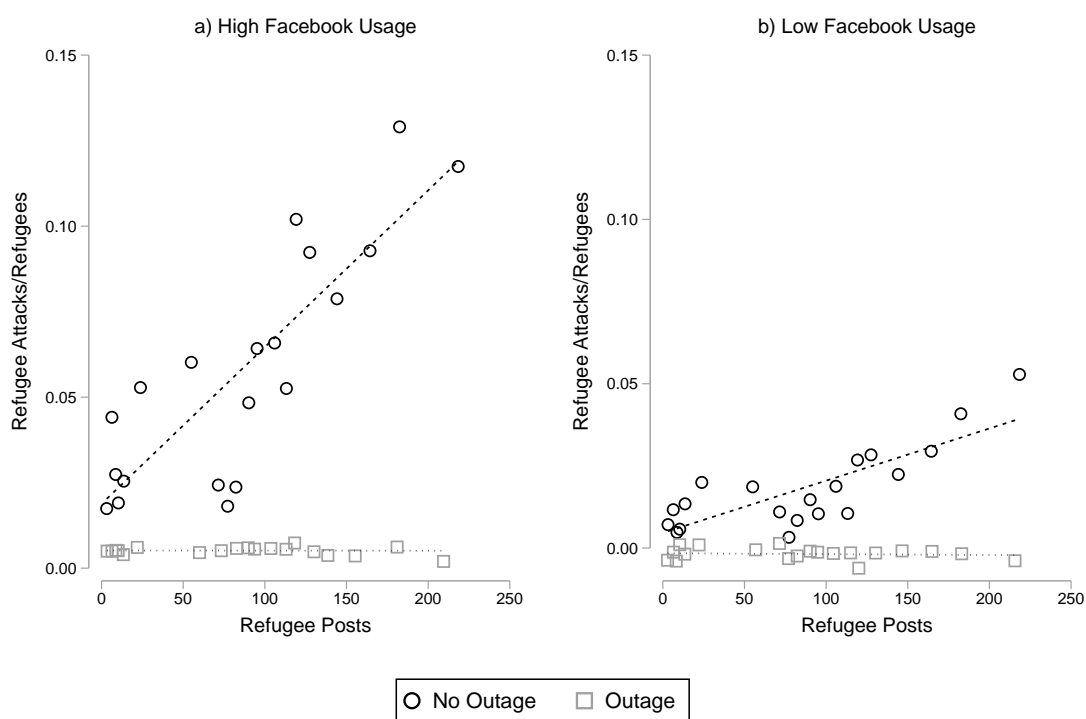
²⁵These posts and comments are the sub-sample of total posts and comments on the AfD page by users who publicly disclosed their location in their Facebook profiles.

Figure 5: Do Local Internet Outages Reduce Local Facebook Activity?



Notes: This figure plots the arithmetic mean and 95% confidence intervals of local Facebook activity measures based on linking users' locations to their posts and comments. The bars marked "Outage" are municipality-week observations in which a local internet outage occurs. The average values are below one, since we do not observe a post or a comment from each municipality in every week. For example, a mean value of around 0.08 for local comments during weeks without outages implies that on average we observe 1 comment for every 12.5 municipality-weeks pairs (out of 480,963) in our data.

Figure 6: Quasi-Experimental Results from Internet Outages



Notes: This figure plots the average number of anti-refugee attacks against our measure of anti-refugee sentiment for municipalities above and below the median of $AfD\ Users/Pop$. Refugee attacks are binned by 20 quantiles of anti-refugee posts, with square-shaped points represent those with outages. The municipalities are additionally split by whether they experience an internet outage in a given week. The number of anti-refugee attacks was residualized with respect to population; hence, the number of attacks can be slightly below 0 in some bins.

To investigate the effects of internet outages more formally, we estimate versions of Equation (2) and report the regression results in Table 5. Column (1) shows that internet outages reduce anti-refugee violence. The coefficient of -0.026 implies that, during such outages, refugee attacks are 70% lower relative to the dependent variable mean (≈ 0.037). Column (2) implies that this effect is driven by periods of high sentiment; it may also be driven by areas with many AfD Facebook users (column 3) but the coefficient is not statistically significant.

In columns (4) through (6), we estimate the full triple-difference model. Here, we estimate the effect of outages in areas with high social media use at times of high anti-refugee sentiment. The estimates suggest that internet problems reduce social media’s impact on anti-refugee violence. While the coefficient of refugee posts and social media exposure is similar to our baseline correlations, the triple interaction term with internet outages is negative and statistically significant in all three specifications. Quantitatively, internet outages appear to mitigate the entire effect of social media. In line with the graphical evidence in Figure 6, we find that the triple interaction coefficient is larger than the baseline coefficient. Put differently, for a given level of anti-refugee sentiment, there are fewer attacks in municipalities with high Facebook usage during an internet outage than in municipalities with low Facebook usage *without* an outage.

Could it be that the effect of internet outages is merely coincidental? As an alternative way of assessing our results’ statistical significance, we perform a randomization test. Instead of the actual internet disruptions, we randomly define 313 municipality-week pairs as placebo outages. We then estimate the same regression using 500 different sets of placebo outages. This allows us to evaluate the probability of finding a statistically significant coefficient in our sample. Using this procedure, we find that 97% of the placebo triple interaction coefficients exhibit lower t -stat. Our findings are thus unlikely to be purely coincidental. We show the full distribution of t -statistics from this randomization test in Figure A.6 in the online appendix.

The identifying assumption for internet outages in our framework is that they only have an effect on anti-refugee hate crime through the reduced exposure to social media. Could it be that we observe reduced hate crimes because users are cut off from the internet generally, not from social media in particular? Two pieces of evidence support the idea we might indeed be capturing social media.

First, when we include interactions of internet disruptions with measures of internet usage (broadband access and per capita internet domains), our main coefficient is unaffected (see column (6) of Table 5). The coefficients of the internet usage interactions are statistically

Table 5: Local Internet Outages and Social Media Transmission

	(1)	(2)	(3)	(4)	(5)	(6)
Baseline Coefficients						
AfD users/Pop. \times Refugee posts				0.369*** (0.113)	0.285*** (0.099)	0.285*** (0.099)
AfD users/Pop. \times Posts \times Outage				-2.037** (0.825)	-2.107** (0.830)	-2.027** (0.811)
Additional Outage Coefficients						
Outage	-0.026*** (0.008)	-0.011 (0.014)	-0.018* (0.010)	-0.029 (0.019)	-0.030* (0.016)	-0.062 (0.077)
Outage \times Refugee posts		-2.278* (1.323)	-	2.467 (1.922)	3.023* (1.768)	2.603 (1.609)
AfD users/Pop. \times Outage			-0.003 (0.003)	0.007** (0.003)	0.006** (0.003)	0.006** (0.003)
Internet Usage Interaction						
Share broadband access \times Outage						0.000 (0.001)
Internet domains/Pop. \times Outage						0.002 (0.001)
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
All other controls (22) \times Posts					Yes	Yes
Observations	480,963	480,963	480,963	480,963	475,302	475,302
Number of municipalities	4,333	4,333	4,333	4,333	4,282	4,282
R^2	0.045	0.045	0.045	0.045	0.046	0.046

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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. Internet outages are defined as municipality-weeks that are in the top quartile of the ratio of reported internet outages to population. Columns (1) through (4) include the baseline controls. Columns (5) and (6) include all controls as in column 7 of table 3, interacted with *Refugee posts* (unreported). Column (6) further adds the interaction of broadband access and internet domains/pop. with local internet outages. Robust standard errors in all specifications are clustered by municipality. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

insignificant. This is at least some indication that we are not merely capturing general internet usage.²⁶ It also suggests that our findings are unlikely to capture that people are busy fixing internet access problems. If we were merely capturing such displacement effects, one would expect it to more strongly affect people in areas with high internet usage, which does not seem to be the case in the data.

Second, internet outages do not attract a consistent statistically significant coefficient after including the other interaction terms in columns (4) through (6). The triple interaction with Facebook usage remains highly significant. This result also supports the idea that internet outages reduce hate crime by limiting access to social media.

Another concern could be that hate crimes are less likely to be reported during internet outages. We believe this is unlikely to explain our findings because we analyze incidents that happened month if not years in the past. While internet outages might hamper the flow of information, it seems highly unlikely that incidents such as assault or arson are *never* reported due to a temporary internet disruption. This should be especially true for official reports by the police. We therefore repeat our analysis for the more than 50% of total anti-refugee incidents in our sample reported by official sources (German Parliament or police). Column (1) of Table A.11 shows that our results also hold in this sample.

The outage results are also robust to using a leave-own-area-out measure of refugee posts (see column 2).²⁷ We can also replicate our finding using a linear probability model with a dummy variable for anti-refugee incidents as an outcome (see column 3). This alternative functional form increases the statistical significance of the triple interaction.

Facebook outages. As further evidence for the social media transmission mechanism, we use eight major Germany-wide Facebook outages as a source of exogenous variation specific to social media access. Table A.4 outlines the details of each of the eight outages and links to relevant press reports. The advantage of these outages is that they are Facebook-specific and therefore do not affect other potential channels of online hate speech transmission. In fact, we find that these outages are uncorrelated with the total number of weekly internet disruptions in the time series of 111 weeks ($t = -0.28$, unreported).

We present the results of interacting these variables analogous to the internet outages in Table 6. The results again reveal a clear pattern. The coefficient of -0.015 in column (1) shows

²⁶Indeed, the ratio of AfD Facebook users to population is, if anything, *negatively* correlated with our measure of internet affinity based on broadband access.

²⁷Note that the increase in the coefficient size we find here is because the post variable based on the geotagged subsample has a lower average number of posts.

that anti-refugee incidents are around 40% lower in weeks with major Facebook outages (relative to the mean number of attacks, which is around 0.38). Column (2) shows that this effect is larger in areas with many users on the AfD Facebook page.²⁸ This is additional evidence that social media *per se* might affect hate crimes.

Next, we introduce the triple interaction of Facebook outages with social media usage and our refugee salience measure. The triple interaction is negative and statistically significant in all three specifications in columns (3) through (5). Quantitatively, we find that Facebook disruptions *fully* undo the baseline correlation of refugee attacks and exposure to social media sentiment. For example, consider that the coefficient of *AfD users/Pop.* and *Refugee Posts* is 0.398 in column (4) but -0.327 on the triple interaction. This implies that, in weeks of major Facebook outages, heightened anti-refugee sentiment is essentially not associated with a differential increase of anti-refugee attacks in municipalities with higher Facebook usage. The interaction of Facebook outages with AfD users is now also insignificant and close to zero. One reading of this is that outages matter particularly at times of high anti-refugee sentiment.

Note that we would expect the Facebook outage coefficients to differ in magnitude from the internet outage coefficients. This is because Facebook outages eliminate the differential exposure *between* areas with high and low social media usage to anti-refugee posts. In contrast, internet outages further exploit variation *within* municipalities. Because within-municipality variation induced by internet outages appears to matter more in our setting, we find smaller coefficients for Facebook outages.

We again perform a randomization test to assess the statistical significance of the Facebook outage results. We randomly assign placebo Facebook outages to eight weeks in our data, excluding the weeks in which we identified Facebook outages. We then estimate the same regression using 500 different sets of placebo outages. Using this procedure, we find that 95% of the placebo triple interaction coefficients exhibit smaller *t*-stats. We show the full distribution of *t*-statistics from this randomization test in Figure A.7 in the online appendix. This confirms that our findings are unlikely to be a matter of coincidence.

Taken together, the evidence here suggests that the relationship of anti-refugee sentiments online and hate crimes is attenuated by Facebook and internet outages. These results are most consistent with a causal propagation effect of social media. They are unlikely to be driven by unobserved municipality characteristics.

²⁸In unreported results, we also find that the interaction of Facebook outages with refugee posts has a statistically significant negative coefficient.

Table 6: Facebook Outages and Social Media Transmission

	(1)	(2)	(3)	(4)	(5)
Baseline Coefficients					
AfD users/Pop. \times Refugee posts			0.398*** (0.121)	0.320*** (0.107)	0.320*** (0.107)
AfD users/Pop. \times Posts \times Outage			-0.327* (0.185)	-0.386** (0.189)	-0.386** (0.189)
Additional Outage Coefficients					
Outage	-0.015*** (0.004)	-	-	-	-
AfD users/Pop. \times Outage		-0.003** (0.001)	0.000 (0.001)	0.000 (0.001)	0.003 (0.002)
Week FE		Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes
Baseline controls (7) \times Posts	Yes	Yes	Yes	Yes	Yes
All other controls (22) \times Posts				Yes	Yes
All controls (29) \times Outages					Yes
Observations	484,293	480,963	480,963	475,302	475,302
Number of municipalities	4,333	4,333	4,333	4,282	4,282
R^2	0.045	0.045	0.045	0.046	0.046

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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. Facebook outages refer to weeks in which Facebook experienced considerable disruptions; see the online appendix for more details on how these are defined. Note that the other interaction terms *Outage*, *Refugee posts* and *Outage \times Refugee posts* are absorbed by the week fixed effects in columns (3) through (5). Columns (1) through (3) include the baseline controls. Columns (4) and (5) include all controls as in column 7 of table 3, interacted with *Refugee posts*. Column (5) adds the interaction of these control variables with Facebook outages. Robust standard errors in all specifications are clustered by municipality. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

3.5 Additional Results

3.5.1 Alternative Measures of Social Media Usage

Our baseline measure of local Facebook penetration is the number of AfD page users divided by population. While this proxy captures the intuition that AfD users are likely to be most exposed to anti-refugee salience on the AfD page, our results also hold for measures of general Facebook usage. Areas that generally use more Facebook should also be more exposed to social media content, including anti-refugee posts. As explained in Section 2, these general measures are based on the locations of users of the Nutella Germany Facebook page (*Nutella Users/Pop.*, and $I_{Many\ Nutella\ Users}$), one of the most popular pages in Germany.

The advantage of these measures is that they are based on an entirely different Facebook page. Further, the $I_{Many\ Nutella\ Users}$ dummy is only correlated with municipality characteristics that one would expect if our measure accurately reflects local differences in social media usage (see Table A.6 in the online appendix for a more detailed discussion). This indicates that we are indeed capturing local Facebook use patterns.

In Table A.8 in the online appendix we replicate the regression results for Equation (1) using these alternative measures of municipal-level Facebook usage. For convenience, the table also reports the baseline results in Panel A. In Panel B the coefficients on the interaction term of *Nutella Users/Pop.* are very similar to our baseline results and still highly statistically significant. In Panel C we use the top tercile of *Nutella users/Pop.* within a county to proxy for general social media usage. We find comparable estimates and the inclusion of controls again has little effect. Table A.13 further replicates our findings on internet and Facebook outages which, again, are largely unchanged using the Nutella proxies. These findings demonstrate that the effects of exogenous reductions in exposure to anti-refugee sentiment do not hinge on using right-wing users to measure local Facebook penetration.

3.5.2 Placebo Tests for other Posts on the AfD Facebook page

If the channel we uncover is indeed specific to refugees, we would expect a weaker correlation between refugee attacks and posts about 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 compare coefficient 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 all these terms, we find no significant relationship between the number of posts 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 correlation we capture 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: Other Facebook Posts and Anti-Refugee Hate Crimes

	Refugee (Baseline) (1)	All other posts (2)	Muslim posts (3)	Islam posts (4)	Jew posts (5)	EU posts (6)
AfD users/Pop. \times 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) \times 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 the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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, with the baseline being *refugee* (“Flüchtling”). Standardized coefficients are 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 levels, respectively.

3.5.3 Results for the Intensive Margin of Facebook Usage

If social media works as the propagating mechanism for hate speech, we would also expect the coefficient magnitudes to increase with AfD Facebook page’s local usage intensity. In the case of Facebook, an intuitive proxy for the time people spend on Facebook is the frequency with which 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 matters particularly 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 the total number of local posts on the AfD wall and the number of comments and likes on AfD posts, all scaled over the number of AfD users in a municipality.²⁹ Note that these user intensity measures are not systematically correlated with local Facebook penetration, city size, or population density. As such, they create additional municipal-level variation.

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, a higher reach per AfD user increases the correlation of social media exposure with hate crimes. Importantly, these interactions 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 coefficient 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 coefficient by around one third.³⁰ 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 estimate is outweighed by the positive coefficient of the triple interaction in any week with more than three refugee posts (the median being 83).

3.5.4 Distracting News Events and Social Media

As an additional piece of analysis, we investigate the role of news shocks on the transmission of online hate speech to real-world actions. Durante and Zhuravskaya (2018) provide evidence 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 topic of refugees. This is somewhat analogous to Facebook outages in that we exploit additional exogenous weekly variation, albeit with a different goal: if we

²⁹Note that we can only construct these measures on the intensive margin of municipalities where we can identify at least one AfD user. Our baseline results also hold in this sub-sample, which we show in Table A.15 in the online appendix.

³⁰To see this, consider that the total implied estimate including interaction is calculated as $0.317 + 0.013 \times 12 \approx 0.47$, which is about 1/3 larger than the baseline coefficient 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(Reach,Population)	0.012	0.012	0.010
Corr(Reach,Pop. density)	0.025	0.038**	0.026
Corr(Reach,AfD users/Pop.)	0.016	0.005	0.019
Corr(Reach,Average age)	-0.019*	-0.011	-0.021*

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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 levels, respectively.

are correct that major news events act as a distraction, they should reduce the correlation of refugee salience with hate crimes.

To measure these news shocks, we obtain Google Trends data on 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 investigate if elevated anti-refugee salience has a weaker link with hate crime in weeks with major news events. If this is the case, we would expect that these events *decrease* the correlation of social media transmission with refugee attacks.

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, changes in anti-refugee incidents correlate less with heightened refugee salience. As the salience of other events crowds that of refugees, there are smaller increases of hate crimes in municipalities with more AfD social media users.

3.6 Additional Robustness Exercises

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

First, we consider a dynamic specification with a lagged dependent variable in column (1) of Table A.14.³¹ 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 coefficient, almost equivalent to the baseline result of 0.369. In column (2), we weight the regression by municipality population to check the possibility that minor villages or population density are driving the result. However, we find the result is hardly affected by this perturbation. The same holds true for column (3), where we

³¹Note that Nickell bias is unlikely to play a role in this dynamic specification because we have a relatively large time series of 111 observations per geographical unit. This is also supported by the fact that including a lag of the dependent variable makes little difference to the point estimate.

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

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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% levels, respectively.

replace the total number of AfD users with the pre-sample number of users to address potential endogeneity concerns that the page’s popularity results from increased refugee violence.³²

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 anti-refugee attacks. The interaction term is highly significant with a t -stat of 3.27. Note that, because the posts are now based on a geotagged subsample, the coefficient size increases due to the lower average number of posts. 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 commenting on existing posts might be different from creating new ones. In column (7), we account for popularity by weighting the number of posts by the number of likes they receive. Finally, in column (8) we use the *share* of anti-refugee posts on the AfD Facebook page in a given week. The results are markedly similar independent of the salience measure.

In Table A.15 in the online appendix, we account for the skewed distribution our AfD Facebook use measure. As shown in Figure A.8, some municipalities have an unusually high ratio of AfD users per capita. To determine whether 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. Our results also hold when we exclude municipalities below the 10th and above the 90th percentile in AfD Facebook usage. Moreover, we show that the correlation of local social media exposure with hate crime 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 one AfD user and interact them separately with the refugee salience trend. The first quartile serves as the excluded group. Strikingly, the results in column (5) suggest that the coefficient’s magnitude of Germany-wide anti-refugee sentiment increases monotonically with the number of local AfD page Facebook users.

We also assess our results’ robustness to different transformations of the refugee attack

³²As explained above, using the overall number of users in a municipality is preferable because it gives us more data on user locations.

variable and estimation methods in Table A.16. 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 $\text{Log}(1 + \textit{Refugee attacks})$ on the interaction of refugee posts and $\text{Log}(1 + \textit{AfD Users})$. This unscaled specification controls for population differences through the interaction of our salience measure with the number of inhabitants as a control variable. The coefficient on the interaction of interest is again highly statistically significant. A log-log version of the main specification Equation (1) also yields similar results (column (6)).

In Table A.17, we also show a range of different standard errors for our baseline findings. As it turns out, clustering by municipality is, overall, a highly conservative choice. Finally, our results are robust to different levels of aggregation. In online appendix Table A.18 we show that the findings are similar if instead at the municipality level we aggregate our data at the county or state level.

4 Conclusion

Social media has become a powerful tool for sharing and disseminating information. In this paper, we investigate whether social media 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 light on how online posts are correlated with anti-refugee incidents in Germany. Plausibly exogenous variation in disruptions to users' Facebook or internet access supports the view that some of this correlation reflects a causal effect.

Existing research shows local cultural attitudes towards foreigners are enormously persistent (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 can be blurred, regulation can open the door to censorship. Our work does, however, suggest that policymakers ignore online hate speech at their peril. Future research should investigate effective ways to tackle online hate speech. By quantifying the extent of the problem, our paper takes a first step towards identifying potential harm arising from extended social media usage.

References

- Adena, M., Enikolopov, R., Petrova, M., Santarosa, V., and Zhuravskaya, E. (2015). Radio and the Rise of The Nazis in Prewar Germany. *The Quarterly Journal of Economics*, 130(4):1885–1939.
- Alesina, A. and La Ferrara, E. (2005). Ethnic Diversity and Economic Performance. *Journal of Economic Literature*, 43:721–761.
- BAMF (2016). Aktuelle Zahlen zu Asyl. *Bundesamt für Migration und Flüchtlinge*.
- Barberá, P. (2014). How social media reduces mass political polarization. evidence from germany, spain, and the us.
- Bartik, T. J. (1991). Boon or Boondoggle? The Debate over State and Local Economic Development Policies.
- BBC (2017). Social Media Warned to Crack Down on Hate Speech.
- Becker, S. O. and Pascali, L. (2016). Religion, Division of Labor and Conflict: Anti-Semitism in German Regions over 600 Years.
- Becker, S. O., Pfaff, S., and Rubin, J. (2016). Causes and Consequences of the Protestant Reformation. *Explorations in Economic History*, 62:1–25.
- Bessi, A., Zollo, F., Del Vicario, M., Scala, A., Caldarelli, G., and Quattrociocchi, W. (2015). Trend of Narratives in the Age of Misinformation. *PLOS ONE*, 10(8):1–16.
- Bhuller, M., Havnes, T., Leuven, E., and Mogstad, M. (2013). Broadband Internet: An Information Superhighway to Sex Crime? *Review of Economic Studies*, 80(4):1237–1266.
- Boxell, L., Gentzkow, M., and Shapiro, J. M. (2017). 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.
- Bursztyn, L., Cantoni, D., Funk, P., and Yuchtman, N. (2017). Polls, the press, and political participation: the effects of anticipated election closeness on voter turnout. *National Bureau of Economic Research*.

- Card, D. and Dahl, G. B. (2011). Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior. *The Quarterly Journal of Economics*, 126(1):103–143.
- Colussi, T., Ispording, I. E., and Pestel, N. (2016). Minority Salience and Political Extremism.
- Dahl, G. and DellaVigna, S. (2009). Does Movie Violence Increase Violent Crime? *The Quarterly Journal of Economics*, pages 677–734.
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., and Quattrociocchi, W. (2016). The Spreading of Misinformation Online. *Proceedings of the National Academy of Sciences*, 113(3):554–559.
- DellaVigna, S., Enikolopov, R., Mironova, V., Petrova, M., and Zhuravskaya, E. (2014). Cross-Border Media and Nationalism: Evidence from Serbian Radio in Croatia. *American Economic Journal: Applied Economics*, 6(3):103–32.
- DellaVigna, S. and Gentzkow, M. (2010). Persuasion: Empirical Evidence. *Annual Review of Economics*, 2(1):643–669.
- DellaVigna, S. and La Ferrara, E. (2015). Economic and Social Impacts of the Media.
- Destatis (2017). *Statistisches Jahrbuch: Deutschland und Internationales*. Statistisches Bundesamt (Destatis), Wiesbaden.
- Durante, R. and Zhuravskaya, E. (2018). Attack when the world is not watching? us news and the israeli-palestinian conflict. *Journal of Political Economy*, 126(3):1085–1133.
- Eisensee, T. and Strömberg, D. (2007). News Droughts, News Floods, and U. S. Disaster Relief. *The Quarterly Journal of Economics*, 122(2):693–728.
- Enikolopov, R., Makarin, A., and Petrova, M. (2016). Social Media and Protest Participation: Evidence from Russia.
- Fernndez-Val, I. and Weidner, M. (2016). Individual and Time effects in Nonlinear Panel models with Large N, T. *Journal of Econometrics*, 192(1):291 – 312.
- Financial Times (2017). Powerhouse Germany Badly Trailing Rivals in Broadband.
- Fiorina, M. P. and Abrams, S. J. (2008). Political Polarization in the American Public. *Annual Review of Political Science*, 11:563–588.

- Focus (2014). Die beliebtesten Fanpages: Adidas, BMW, FC Bayern: Diese Facebook-Seiten lieben Ihre Freunde. *www.focus.de/digital/internet/adidas-bmw-fc-bayern-das-sind-die-10-erfolgreichsten-deutschen-facebook-seiten_id_3967395.html*.
- Fouka, V. and Voth, H.-J. (2013). Reprisals Remembered: German-Greek Conflict and Car Sales during the Euro Crisis.
- Gabler, N. (2016). The Internet and social media are increasingly divisive and undermining of democracy. *Alternet*.
- Gavazza, A., Nardotto, M., and Valletti, T. M. (2018). Internet and Politics: Evidence from U.K. Local Elections and Local Government Policies. *The Review of Economic Studies*, page rdy028.
- Gentzkow, M. (2006). Television and Voter Turnout. *The Quarterly Journal of Economics*, 121(3):931–972.
- Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2018). Bartik Instruments: What, When, Why, and How. *Working Paper*.
- Hölig, S. and Hasebrink, U. (2016). *Reuters Institute Digital News Survey 2017: Ergebnisse für Deutschland*, volume Nr. 38 of *Arbeitspapiere des Hans-Bredow-Instituts*. Verlag Hans-Bredow-Institut, Hamburg.
- Hölig, S. and Hasebrink, U. (2017). *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.
- Jha, S. (2013). Trade, Institutions, and Ethnic Tolerance: Evidence from South Asia. *American Political Science Review*, 107(4):806–832.
- Manacorda, M., Tesei, A., et al. (2017). Liberation Technology: Mobile Phones and Political Mobilisation in Africa. Technical report, Centre for Economic Performance, LSE.
- Müller, K. and Schwarz, C. (2018). Making America Hate Again? Twitter and Hate Crime Under Trump. *Available at SSRN: <https://ssrn.com/abstract=3149103>*.
- New York Times (2016). How Facebook Warps Our Worlds, By Frank Bruni.

- New York Times (2017a). How Fiction Becomes Fact on Social Media, By Benedict Carey.
- New York Times (2017b). Seeking Asylum in Germany, and Finding Hatred, By Ainara Tiefenthaler, Shane Oneill and Andrew Michael Ellis .
- OECD (2016). Broadband Statistics.
- Oksanen, A., Hawdon, J., Holkeri, E., Näsi, M., and Räsänen, P. (2014). Exposure to Online Hate Among Young Social Media Users. *Soul of Society: A Focus on the Lives of Children & Youth*, 18:253–273.
- Pariser, E. (2011). *The Filter Bubble: What the Internet is hiding from you*. Penguin UK.
- Pew Research Center (2016). Social Media Update 2016. Technical report.
- Pew Research Center (2018). News Use Across Social Media Platforms 2018. Technical report.
- Schmidt, A. L., Zollo, F., Del Vicario, M., Bessi, A., Scala, A., Caldarelli, G., Stanley, H. E., and Quattrociocchi, W. (2017). Anatomy of News Consumption on Facebook. *Proceedings of the National Academy of Sciences*, 114(12):3035–3039.
- Stephens-Davidowitz, S. (2014). The Cost of Racial Animus on a Black Candidate: Evidence using Google Search Data. *Journal of Public Economics*, 118:26–40.
- Sunstein, C. R. (2009). *Republic. com 2.0*. Princeton University Press.
- Sunstein, C. R. (2017). *# Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- The Guardian (2017). CPS to Crack Down on Social Media Hate Crime, says Alison Saunders, by Vikram Dodd.
- Voigtlander, N. and Voth, H.-J. (2012). Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany. *The Quarterly Journal of Economics*, 127(3):1339–1392.
- Voigtlander, N. and Voth, H.-J. (2015). 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.
- Yanagizawa-Drott, D. (2014). Propaganda and Conflict: Evidence from the Rwandan Genocide. *The Quarterly Journal of Economics*, 129(4):1947–1994.

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 Greece and Spain received 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 attacking 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. As the Euro Crisis cooled, the party began to shift its focus further to the right on topics like traditional family values or the role of Islam in Germany. These more nationalist-conservative political positions, championed by Frauke Petry, attracted a significant share of far-right recruits to the party. In 2015, Petry was elected the main speaker of the party, a major defeat for its founder, Bernd Lucke. As a result of this loss, Lucke resigned from his leadership position and left the party completely, followed by several other key party members.

In the run-up to the 2017 federal election, the AfD leadership included Frauke Petry, Alexander Gauland, Björn Höcke, Jörg Meuthen, and Beatrix von Storch, all of whom hold staunch national conservative opinions. With the beginning of the refugee crisis, the aggressively framed mass immigration as dangerous and declared they were unwilling to accept any refugees into Germany. This messaging was accompanied by increased xenophobia and criticism of Islam.

Under the new leadership and impelled by the refugee crisis, the AfD continued to win elections, securing seats in 14 out of the 16 state parliaments in 2016. 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

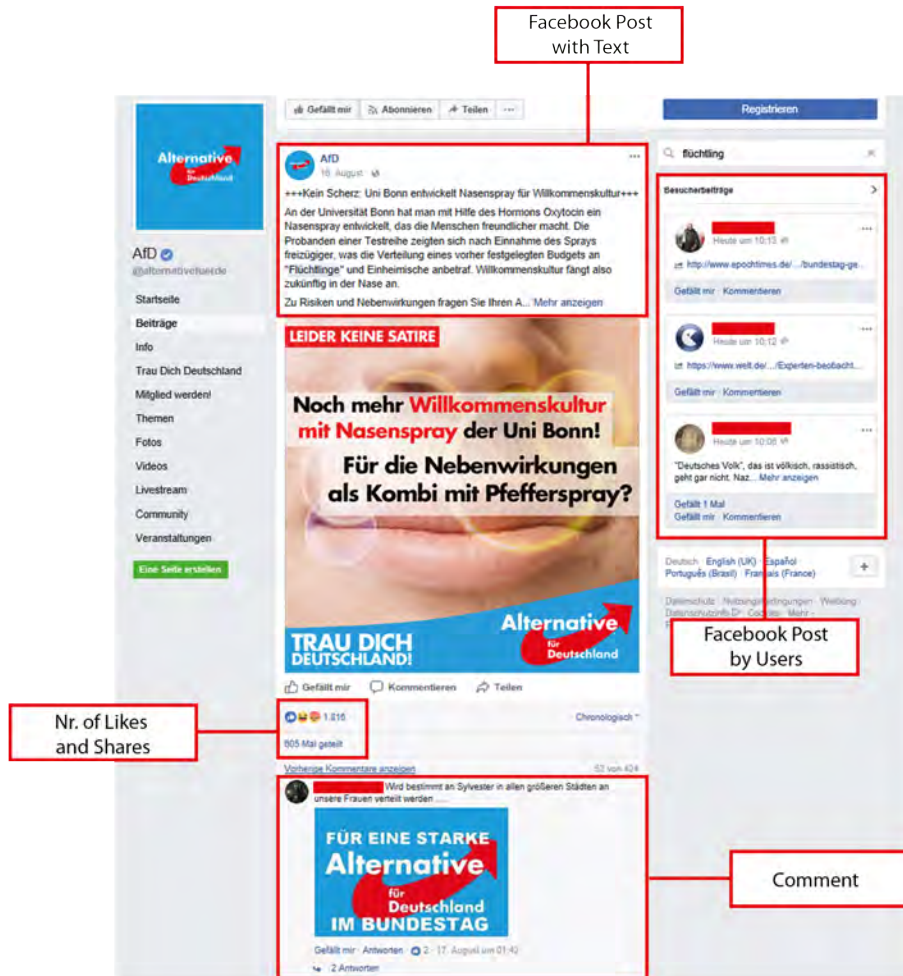
On Facebook, celebrities, universities, restaurants, and political groups like the AfD have created their own Facebook pages. The AfD page is the starting point for its followers on Facebook. Any Facebook user who is interested in or supports the AfD can “like” its page. The messages posted on the AfD’s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

In addition to receiving information from the AfD, Facebook users can become active on the party’s page as well. In general, such interactions fall into three categories. First, people can post their own messages, links, or pictures on the fan page. These posts are visible to everybody but will not automatically appear in other users’ Facebook feeds. Second, users can comment on posts and comments by other users or the AfD itself. Those comments appear below the original post and are also visible to the public. Third, each post or comment can be “liked” as a sign of support.

Figure A.1 shows an example of how these three interaction types show up on the AfD page. The Facebook Graph API allowed us to collect all post, comments, and likes from the AfD’s fan page, information we highlight in Figure A.1. Facebook assigns each user a unique ID that makes it possible to attribute posts and comments to individual profiles.

To hand collect user data, one must visit each individual Facebook user profile, from which, depending on the user’s privacy settings, one can determine his/her 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 publicly available user information extracted from Facebook. The authors removed users' personal information for privacy.

Table A.1: Examples of Anti-Refugee Incidents

Date	Place	Description	Type
03.11.2016	Braunsbedra	A 20-year old Syrian was riding his bike in the evening and approached a man he assumed needed his help. Suddenly, two additional men appeared and all three started kicking and punching the victim.	Assault
28.12.2016	Langenhagen	An unknown person sprayed graffiti on a refugee home. The graffiti read “Deutsch Nantional (German-National, misspelled in original), “18 (code for Adolf Hitler) and “88 (code for Heil Hitler).	Property Damage
17.11.2016	Oschersleben	A fire occurred in a villa which had until recently accommodated refugees. After a forensic analysis, the police concluded it to be a case of arson, since the fire started in several places at once using fire accelerant. Furthermore, a detonation occurred when the police arrived. Nobody was injured.	Arson
30.01.2016	Schmölln	450 people participated in a demonstration of the “Thügida (Pegida in Thuringia). The police charged 4 people with violating gun control laws and the Public Meetings Act.	Demonstration
30.01.2016	Berlin	The police investigated an insult against inhabitants of a refugee home.	Suspected Cases

Notes: This table reports one example for each class of anti-refugee incidents in the data. The descriptions were translated from German into English by the authors.

Table A.2: Translated Example AfD Posts from Facebook

Date	Post	# Likes	# Comments	# Shares
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 five 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.

B.2. Validation of Internet and Facebook Outages

Figure A.3: Map of Internet Outages



Notes: This map plots the geographic distribution of internet outages per million inhabitants for the German municipalities in the data. These outages are defined as the total number of user-reported disruptions retrieved from *Heise.de*. See Section 2 for more details.

Table A.4: Validation of Facebook Outage Data

Peak Date	Description	Source
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 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 Germany, the Netherlands, and the United Kingdom.	Link
15/07/2015	Facebook suffered a 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 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 particularly affected.	Link
16/06/2016	Facebook had an outage concentrated in Western Europe. Users were unable to log in, post, use the messenger, or could not access 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, post, or 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.

Table A.3: Validation of Internet Outage Data
(a) Part 1/2

Date	Provider	Region	Description	# Outages	Source
12/06/2015	Kabel D. and Unitym.	Germany	The IT website "Netzwelt" reported 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	On July, 18th 2015, the news site "KA News" reported a disruption at internet provider Kabel BW, a subsidiary of Unitymedia. Kabel BW 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 disrupted internet connections 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 1u1 reported disruptions of their internet and phone connections. The problems had started on the June 27th and were largely fixed by the evening of the 5th. Neither provider explained what had caused the problems.	27	Link
08/07/2015	Versatel	Münster	The "Halteiner Zeitung" reported on July 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 Münster area were also affected by the problem.	29	Link
20/08/2015	Unitymedia	NRW and Hessen	The "Gießener Allgemeine" reported on August, 20th 2015 that many users of internet provider Unitymedia encountered disrupted internet connections beginning 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

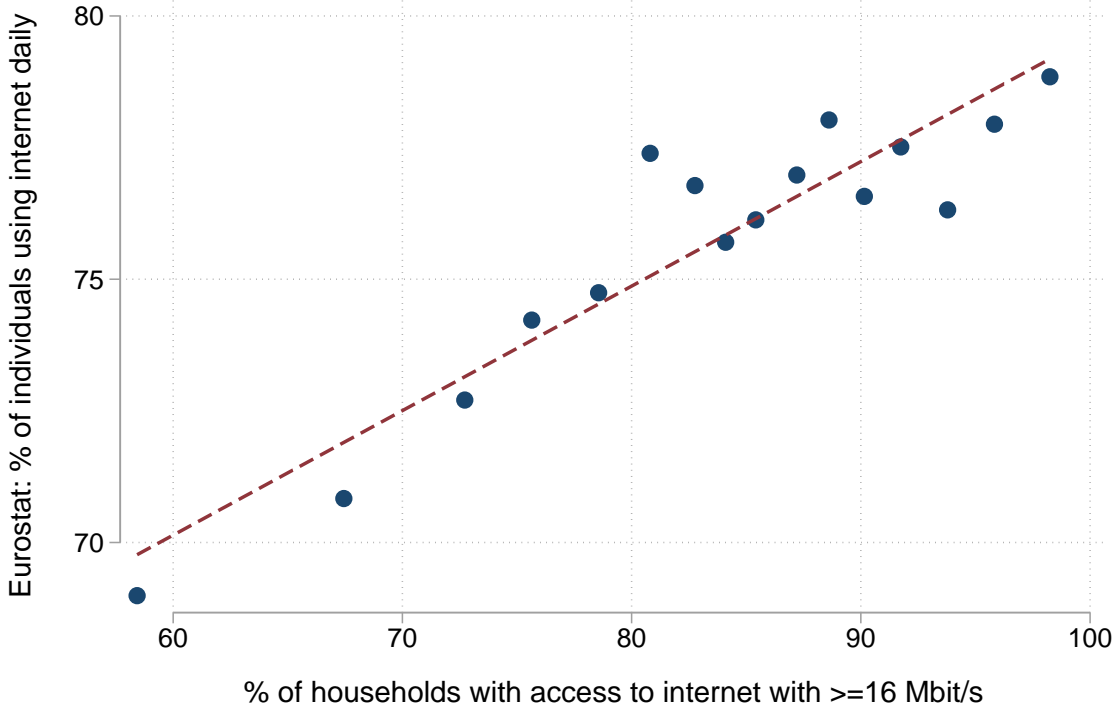
(b) Part 2/2

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 disrupted internet connections beginning in the early morning of the same day. Most of the reports came from the major cities Berlin, Hamburg, Munich, and Frankfurt. According to the German Telekom, the problem was caused by a breakdown of a RADIUS server that is responsible for authenticating 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 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 with the provider T-Online. The outage affected not only 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 a Germany-wide outage of the internet provider O2. The problems were concentrated in metropolitan areas. At the time of report O2 was still investigating the cause of the problem, which was likely an issue with the company’s VoIP system.	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 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 on 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 a link to the original news source. The column “# Outage” refers to the number of outages of the affected provider reported by users on Heise.de, which serve as the basis of our internet outage measures; note that this number reflects the number of user reports, *not* the actual number of affected users. The web pages were last accessed in February 2018.

B.3. Additional Variable Overview

Figure A.4: Daily Internet Users and Share of Households with Broadband Access



Notes: This figure plots the municipal-level share of households with access to broadband internet (≥ 16 Mbit/s) against the state-level percentage of individuals using the internet daily taken from Eurostat survey data, binned into 16 quantiles. The corresponding correlation coefficient is 0.9245.

Table A.5: Overview Variables

(a) Part 1/4

Variable	Level	Description	Source
Refugee Attacks			
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
Social Media Data			
AfD Users/Pop.	Municipality	The number of AfD Users in each municipality divided by population.	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
Auxiliary Variables			
$I_{Internet\ Outage}$	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
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
Baseline Controls			
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
Raw Data			
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
Additional Media and Internet Controls			
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, which we impute using the population, population density, AfD vote share, and county fixed effects (results are almost equivalent without imputation).	ZMG
Additional Right-Wing Controls			
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 who were deported during Nazi times. To analyze 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 anti-Semitic newspaper published by Nazi politician Julius Streicher. To analyze cross-sectional correlates, we scale the number of letters by the population in 1933 (see text for details).	Voigtlander and Voth (2012)
Additional Socio-Economic Controls			
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 manufacturing employees in a given county.	Destatis

(d) Part 4/4

Variable	Level	Description	Source
Additional Voting Controls (2013 & 2017 Election)			
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 "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
Additional Demographic Controls			
Share aged 0-24	Municipality	The number of people aged 0-24, divided by population.	Destatis
Share aged 25-49	Municipality	The number of people aged 25-49, divided by population.	Destatis
Share aged 50-74	Municipality	The number of people aged 50-74, divided by population.	Destatis
Share aged above 75	Municipality	The number of people aged 75 and up, divided by population.	Destatis

In this table, we investigate correlates of the $I_{Many\ Nutella\ Users}$ social media measure used in the paper. More precisely, we report the estimated coefficients from regressing the variables in the left column (one-by-one) on the $I_{Many\ Nutella\ Users}$ measures. All coefficients are standardized to have a mean of 0 and standard deviation of 1 to make magnitudes comparable. 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 are on the municipal-level. $I_{Many\ Nutella\ Users}$ is 1 for municipalities in the top tercile of Nutella users per capita within a county, and 0 otherwise.

As described in the main text, $I_{Many\ Nutella\ Users}$ is only correlated with municipality characteristics that are consistent with general social media usage patterns in Germany. More concretely, we find that younger people 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 Hölzig and Hasebrink (2016), people aged 18-34 are twice as likely to use social media as their main news source; as a result, they are less willing to pay for newspaper subscriptions (Hölzig and Hasebrink, 2017). Finally, we find more Facebook users per capita in urban areas and larger cities, a pattern that has been widely documented in surveys on social media demographics (e.g. Pew Research Center, 2016).

Table A.6: Correlates of $I_{Many\ Nutella\ Users}$

	$I_{Many\ Nutella\ Users}$		$I_{Many\ Nutella\ Users}$	
	$\hat{\beta}$	S.E.	$\hat{\beta}$	S.E.
Baseline Controls				
Population (2015) [†]	0.2236***	(0.0364)	CDU vote share	0.0245 (0.0305)
GDP/Worker	0.0481	(0.0305)	SPD vote share	0.0026 (0.0301)
Population density	0.2360***	(0.0323)	Left vote share	-0.0271 (0.0304)
AfD vote share (2017)	-0.0467	(0.0301)	Green vote share	0.0352 (0.0304)
Share Abitur	-0.0347	(0.0303)	FDP vote share	-0.0072 (0.0302)
Share broadband access	0.0497	(0.0307)	Pirate vote share	0.0414 (0.0304)
Share immigrants	0.0482	(0.0315)	Voter turnout	0.0028 (0.0304)
Additional Media and Internet Controls[†]				
Internet outages/Pop.	0.0333	(0.0339)	Additional Socio-Economic Controls	
Registered domains/Pop.	0.0472	(0.0312)	Average age	-0.0541* (0.0308)
News paper sales/Pop.	-0.1390***	(0.0311)	Share benefit recipients	-0.0181 (0.0305)
Additional Right-Wing Controls				
Nazi murders/Pop. [†]	0.0292	(0.0294)	Share non-Christians	-0.0266 (0.0308)
NPD vote share (2017)	-0.0643**	(0.0301)	Manufacturing share (%)	-0.0309 (0.0303)
Deported Jews/Jews (1933)	0.0007	(0.0671)	Additional Demographic Controls	
Stuermer letters/Pop. (1933)	-0.0345	(0.0619)	Share aged 0-24	0.0388 (0.0300)
			Share aged 25-49	0.0584* (0.0299)
			Share aged 50-74	-0.0025 (0.0299)
			Share aged 75 and older	0.0602** (0.0301)

Notes: This table reports estimated coefficients from regressing the variables in the left column (one-by-one) on the $I_{Many\ Nutella\ Users}$ measure. All coefficients are standardized to have a mean of 0 and standard deviation of 1 to make magnitudes comparable. 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 municipal-level. The $I_{Many\ Nutella\ Users}$ dummy is 1 for municipalities in the top tercile of Nutella users per capita within a county, and 0 otherwise. See text for discussion.

C Additional Results: AfD Activity and Anti-Refugee Incidents

Table A.7: Social Media and Hate Crime: Time Series Correlations

Dependent variable:	Number Refugee attacks		
	All municipalities (1)	AfD users > 0 (2)	AfD users = 0 (3)
Anti-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 anti-refugee attacks on the number of anti-refugee posts on the AfD's Facebook page. We standardize the number of anti-refugee attacks to make the coefficients comparable across the different samples. The regression is reported for the full sample (column 1), for municipalities with at least one AfD Facebook user (column 2), and for municipalities with no AfD Facebook users (column 3). Robust standard errors are in parentheses. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

Figure A.5: Nutella Facebook Usage per Capita, by Municipality



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 Nutella Germany.

Table A.8: Baseline Results with Alternative Social Media Measures

	Additional interacted controls						
	Baseline controls (1)	Right-wing controls (2)	Media controls (3)	Socio-economic controls (4)	2017 vote controls (5)	Age structure controls (6)	All controls (7)
Panel A: AfD Users/Pop.							
Social media users/Pop. \times Refugee posts	0.369*** (0.113)	0.360*** (0.111)	0.325*** (0.105)	0.355*** (0.112)	0.346*** (0.108)	0.346*** (0.113)	0.285*** (0.098)
Observations	480,963	480,963	480,963	475,302	480,963	480,963	475,302
R^2	0.045	0.045	0.045	0.046	0.045	0.045	0.046
Panel B: Nutella Users/Pop.							
Social media users/Pop. \times Refugee posts	0.287*** (0.091)	0.283*** (0.092)	0.233*** (0.087)	0.258*** (0.093)	0.263*** (0.091)	0.252*** (0.091)	0.210** (0.089)
Observations	480,963	480,963	480,963	475,302	480,963	480,963	475,302
R^2	0.045	0.045	0.045	0.046	0.045	0.045	0.046
Panel C: $I_{Many\ Nutella\ Users}$							
Social media users/Pop. \times Refugee posts	2.317*** (0.621)	2.297*** (0.617)	2.206*** (0.606)	2.228*** (0.604)	2.182*** (0.600)	2.210*** (0.620)	1.883*** (0.562)
Observations	480,852	480,852	480,852	475,191	480,852	480,852	475,191
R^2	0.045	0.045	0.045	0.046	0.046	0.045	0.046
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes


Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-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, 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. All control variables are interacted with the $Refugee\ posts$ measure; see text for a description of the controls. Robust standard errors in all specifications are clustered by municipality. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

Table A.9: Baseline 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 the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-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, 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. 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 levels, respectively.

Figure A.6: Randomization Test for Internet Outages



Notes: This figure shows the results of the randomization test, in which we randomly assign placebo internet outages to 313 municipality-week pairs. We repeat this process 500 times and save the t -stat of the triple interaction term of interest. The vertical line marks the t -stat of the actual estimate.

Figure A.7: Randomization Test for Facebook Outages



Notes: This figure shows the results of the randomization test, in which we randomly assign placebo Facebook outages to eight weeks in our data. We repeat this process 500 times and save the t -stat of the triple interaction term of interest. The vertical line marks the t -stat of the actual estimate.

Table A.10: Do Internet Outages Reduce Facebook Activity? More Evidence

	Refugee posts		Refugee likes		Total posts		Total likes		Total comments	
	Full	Posts > 0	Full	Likes > 0	Full	Posts > 0	Full	Likes > 0	Full	Comments > 0
Panel A: 623 Outages, Values Above Median										
No outage	0.008	0.144	0.017	0.383	0.069	0.283	0.153	0.860	0.082	0.161
Internet outage	0.002	0.019	0.000	0.000	0.016	0.034	0.055	0.155	0.072	0.102
Observations	495,726	27,306	495,726	21,756	495,726	120,435	495,726	88,356	495,726	254,634
Diff. > 0?	3.824***	6.337***	19.358***	19.519***	9.426***	17.147***	3.009***	7.261***	0.577	2.349**
Panel B: 313 Outages, Values in Top Quartile (Baseline)										
No outage	0.008	0.144	0.017	0.382	0.069	0.282	0.153	0.859	0.082	0.161
Internet outage	0.003	0.111	0.000	0.000	0.003	0.013	0.000	0.000	0.013	0.028
Observations	495,726	27,306	495,726	21,756	495,726	120,435	495,726	88,356	495,726	254,634
Diff. > 0?	1.466*	0.291	19.358***	19.519***	16.397***	16.845***	26.264***	26.349***	8.580***	7.481***

Notes: This table presents the arithmetic mean of local measures of Facebook activity based on linking user locations with posts, received likes, and comments. 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, or comment over time, respectively. The internet outage dummy in panel A tags positive municipality-week observations that are above the median in the ratio of user-reported internet problems to population; the dummy in panel B is based on the top quartile (the baseline measure used in the regressions).

Table A.11: Robustness: Ruling out Alternative Channels

	(1)	(2)	(3)
	Official Reports	Leave One Out Estimator	Refugee Attack Dummy
Outage	-0.001 (0.056)	-0.047 (0.078)	-0.006 (0.008)
Posts \times Outage	1.581 (1.307)	3.239 (3.084)	0.153 (0.164)
Social media users/Pop. \times Outage	0.005** (0.002)	0.005* (0.003)	0.001 (0.000)
Social media users/Pop. \times Refugee posts	0.156*** (0.060)	0.701*** (0.239)	0.030*** (0.010)
Social media users/Pop. \times Posts \times Outage	-1.612** (0.738)	-4.103*** (1.511)	-0.204*** (0.059)
Week FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Baseline controls (7) \times Posts	Yes	Yes	Yes
All other controls (22) \times Posts	Yes	Yes	Yes
All controls (29) \times Outages	Yes	Yes	Yes
Observations	475,302	475,302	475,302
Number of municipalities	4,282	4,282	4,282
R-squared	0.034	0.046	0.083

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable in Column (1) and (2) is the ratio of municipal-level refugee attacks to the number of asylum seekers. In Column (3) the dependent variable is a dummy variable for refugee attacks. *AfD users/Pop.* is the ratio of people with any activity on the AfD 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) only uses anti-refugee incidents based on official reports (police or parliament), which are unlikely to be subject to time-varying reporting bias. In column (2) we construct a leave one out measure of *Refugee posts*. Internet outages are defined as municipality-weeks that are in the top quartile of the ratio of reported internet outages to population. Columns (1-3) include all controls as in column 7 of table 3, interacted with *Refugee posts* (unreported) and the interaction of the controls with the outage dummy. Robust standard errors in all specifications are clustered by municipality. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

Table A.12: 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 (7) \times Posts	Yes	Yes	Yes

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (2). The dependent variable is the ratio of municipal-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 levels, respectively.

Table A.13: Robustness: Outages and Other Social Media Measures

	Local Internet Outages		Country-Wide Facebook Outages	
	Nutella Users/Pop. (1)	$I_{Many\ Nutella\ Users}$ (2)	Nutella Users/Pop. (3)	$I_{Many\ Nutella\ Users}$ (4)
Outage	-0.019 (0.026)	-0.027 (0.025)	-	-
Outage \times Posts	-0.092 (2.144)	-0.019 (2.065)	-	-
Social media users/Pop. \times Outage	0.005 (0.007)	0.039 (0.027)	0.000 (0.000)	0.001 (0.009)
Social media users/Pop. \times Refugee posts	0.287*** (0.091)	2.321*** (0.622)	0.317*** (0.098)	2.588*** (0.664)
Social media users/Pop. \times Posts \times Outage	-1.108* (0.662)	-5.160** (2.587)	-0.337** (0.172)	-2.991*** (1.049)
Week FE	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes
Baseline controls \times Posts	Yes	Yes	Yes	Yes
Observations	480,963	480,852	480,963	480,963
Number of municipalities	4,333	4,332	4,333	4,332
R^2	0.045	0.045	0.045	0.045

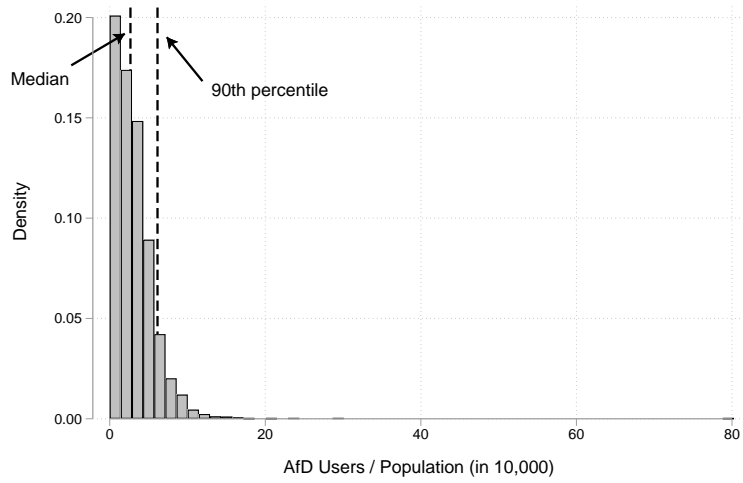
Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (2). The dependent variable is the ratio of municipal-level refugee attacks to the number of asylum seekers. $Nutella\ users/Pop.$ is the ratio of people with any activity on the Nutella Facebook page 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. 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 $Outage \times Posts$ 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 levels, respectively.

Table A.14: Further Robustness Checks

	Refugee Attacks/ Asylum Seekers Lagged Dependent Variable (1)	Pop.- Weighted Regression (2)	Pre- Sample Users (3)	Winsorize Attacks (4)	Leave One Out Estimator (5)	Refugee Comments Interaction (6)	Refugee Likes Interaction (7)	Refugee Post Share (8)
AfD users/Pop. \times Refugee posts	0.368*** (0.112)	0.339*** (0.093)	0.463*** (0.146)	0.282*** (0.086)	0.864*** (0.264)	1.029*** (0.285)	0.002*** (0.001)	0.054*** (0.015)
AfD users/Pop. \times Refugee comments								
AfD users/Pop. \times Refugee post likes								
AfD users/Pop. \times Refugee post share								
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) \times 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 ²	0.046	0.050	0.045	0.052	0.045	0.045	0.045	0.045

Notes: This table presents the estimated coefficients from a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment 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) uses 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 posts on the AfD Facebook 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. All control variables are interacted with the *Refugee posts* measure; see text for a description of the controls. Robust standard errors in all specifications are clustered by municipality. ***, **, * and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

Figure A.8: 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 vertical lines indicate the 50th and 90th percentile of the distribution, respectively, which we make use of in Table A.15.

Table A.15: Accounting for the Skewed Distribution of AfD Users

Dependent variable AfD user percentiles	Anti-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 a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-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 levels, respectively.

Table A.16: Alternative Transformations and Estimation Methods

	Refugee Attack Dummy		Log 1 + Refugee Attacks		Log 1 + Refugee Attacks/ Asylum Seekers	
	OLS (1)	Logit (2)	OLS (4)	OLS (5)	OLS (6)	OLS (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
Municipality FE	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 a regression of hate crimes against refugees on the interaction of local social media usage and anti-refugee sentiment 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 Fernandez-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. All control variables are interacted with the $Refugee\ posts$ measure; see text for a description of the controls. Robust standard errors in all specifications are clustered by municipality. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

Table A.17: Robustness: Main Estimations 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 the interaction of local social media usage and anti-refugee sentiment as in Equation (1). The dependent variable is the ratio of municipal-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 levels, respectively.

Table A.18: 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 municipal-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 levels, respectively.