

Public Signal and Private Action: Right-wing Protest and Hate Crimes

Sulin Sardoschau*

Annalí Casanueva-Artís[†]

February 2025

Abstract

When and why do protests spawn extremist behavior? This paper examines the impact of protest on political radicalization on the far-right, specifically violence against minorities. We study this question in the context of PEGIDA, Germany's most prominent right-wing movement since World War II. Exploiting a difference-in-differences design combined with an instrumental variable and event-study approach, we find that a 20% increase in local protest attendance nearly doubles the incidence of hate crimes. Our evidence shows that large protests serve as a signal of others' preferences that emboldens individuals to commit hate crimes, rather than raising issue salience or facilitating coordination. Large protests set a persistently higher equilibrium level of hate crimes but they do not spark counter-protests on the streets, only resistance on social media. Protest signals diffuse through right-wing social media networks, are amplified by newspaper coverage, and embolden recidivist extremists who typically act alone in public spaces.

Keywords: protest, signal, hate crime, refugees, right-wing

JEL classification: D74, J15, D83, Z10, D72

*Humboldt Universität zu Berlin, School of Business and Economics, and IZA. Email: sulin.sardoschau@hu-berlin.de

[†]Ifo Institute, Junior Research Group on Social Media and Center for Industrial Organization and New Technologies. We would like to thank Davide Cantoni, Ernesto Dal Bó, Jesse Shapiro, and Christian Traxler for valuable feedback, and Merilin Martinson, Jonathan Pardo and Güney Güngör for excellent research assistance. We would also like to thank Karim Soubair for sharing complementary data on PEGIDA networks with us. We appreciate the comments made by participants of the Barcelona Summer Forum on Migration, Berlin Applied Micro Seminar, Research Seminar Hohenheim/Stuttgart, Louis-André Gérard-Varet (LAGV) Public Economics Conference Marseille, 8th IBEO Workshop Alghero, Public Governance Working Group PSL, 7th Monash-Warwick-Zurich Text-as-Data Workshop, POLECONUK Webinar, Silaplana Public Economics Workshop, Center for Economic Research Seminar at ITAM, participants of Junior Applied Microeconomists in Europe Seminar.

1 Introduction

Over the past two decades, the number of protests has nearly doubled worldwide with an increasing share turning violent (ACLED, 2024). Policy proposals aimed at banning or restricting such protests weigh freedom of speech against public safety and social order. A central concern is that protests may incite hatred or encourage acts of violence in pursuit of their political goal. Widespread mobilization can sway elections and policies (Cantoni et al., 2024) but when and why do protests ignite extremist behavior with far-reaching social costs?

In this paper, we study the impact of protests on political violence on the far-right. We do so in the context of Germany’s largest right-wing movement since World War II: the Patriotic Europeans Against the Islamisation of the Occident (PEGIDA). PEGIDA has links to neo-Nazi and fascist groups and entered the German mainstream at the height of the refugee inflow in 2015 as a movement of “concerned citizens.” It drew tens of thousands of participants across the country and inspired offshoots abroad. We examine whether larger PEGIDA protests increased hate crimes against minorities and investigate the underlying mechanisms.

We conceptualize radical behavior as a strategic response to shifting costs, benefits, and preferences, each shaped by large protest events. Individuals weigh the risk of sanctions against the payoff from extremist acts (Becker, 1968). Protests can transform this trade-off through three channels. First, large protests function as a public signal of others’ preferences. This signal can be interpreted as social acceptance of anti-minority sentiment, eroding the perceived stigma, reputational or material cost (probability of detection and punishment) of violence (Bursztyn et al., 2020; Bursztyn & Yang, 2022). It can also shift preference parameters when large crowds validate an ideology and offer a “taste” reward for radical action (Passarelli & Tabellini, 2017). Second, they serve as a coordination device that lowers the cost of organizing violence. Third, they create issue salience, heightening emotions and intensifying attention to the underlying grievance. However, widespread protest success may also signal that the political objective can be achieved through less costly, established democratic processes, reducing the relative returns to political violence (Acemoglu & Robinson, 2000).

Leveraging variation across 84 protest locations and 260 weeks between 2015 and 2020, we estimate a linear probability model with two-way fixed effects and complement this with an instrumental variable strategy and event-study approach. Specifically, we instrument protest size with the interaction of scheduled PEGIDA events and “pleasant” weather (dummy for moderate temperature and minimal rainfall), controlling separately for protest occurrence and weather. Importantly, this allows us to isolate the effect of protest at the *intensive* margin and thus assuages many of the concerns related to weather instruments in general, and instrumented protest more narrowly.¹ The specification includes week fixed effects, municipality-by-month-of-year fixed effects, linear municipality-specific trends, as well as a set of time-varying controls, including lagged protest size and hate crimes, as well as economic, political and crime-related

¹Unlike previous studies that use weather as an instrument for protest *occurrence* alone (Madestam et al., 2013; Beraja et al., 2023; Qin et al., 2024), we focus on the *intensive* margin of protest size and control for both protest and weather conditions (alongside their interactions with economic and political covariates). This approach mitigates concerns documented in the broader literature on rainfall-based identification, such as time-varying confounders or serial correlation in weather, and improves upon simpler designs that may confound protest size with direct weather effects (Sarsons, 2015; Mellon, 2021).

variables – and their separate interactions with protest and weather.²

Our main specification thus estimates the effect of a percentage increase in protest size on the occurrence of a hate crime in the same municipality and week, holding constant local characteristics, the direct impact of weather or protest, and secular (linear) time trends or seasonal differences across municipalities. We find that a 1% increase in protest size raises the probability of hate crimes by approximately 0.5 percentage points. To put this into perspective, a 25% increase in protest size, which corresponds to the average effect of pleasant weather on participation, raises the likelihood of a hate crime by about 11 percentage points, more than doubling the baseline probability of 8.1%.

We probe the robustness of our results with a battery of empirical exercises, addressing potential bias in crime reporting, using alternative instruments, controlling for weather on subsequent days, accounting for spatial correlation, mitigating bias from multiple or staggered treatment in two-way fixed effects models, changing sample composition, as well as using alternative definitions for treatment and outcome and more.³

Our estimate captures the local average treatment effect for compliers, meaning protest locations where turnout increases due to pleasant weather. These additional participants are likely marginal supporters rather than the movement’s dedicated base, whose impact is captured by the protest dummy. Importantly, this does not imply that marginal supporters commit these crimes. In fact, as we will show, hate crimes are carried out by already radicalized individuals. However, greater protest turnout may facilitate and encourage hate crimes in different ways.

We identify three potential mechanisms: salience, coordination, and signaling. First, larger protests may heighten the salience of anti-refugee sentiment, reinforcing existing resentment and prompting a temporary spike in hate crimes as individuals are reminded of the issue. Second, protests may serve as coordination hubs, facilitating the collective planning and execution of hate crimes. Third, large protests act as a public signal of widespread anti-refugee sentiment, potentially lowering the perceived social or actual costs of committing hate crimes. In the extreme, this signal may also change individuals’ preferences by increasing the perceived social rewards or by creating “taste” for such actions. We disentangle these mechanisms through a series of empirical tests.

First, we employ an event-study design at the municipality-day and municipality-week level to analyze the dynamic effects of PEGIDA protest on pleasant days on hate crimes. The results reveal no pre-existing differences in hate crime trajectories and show a sharp increase in crimes on pleasant protest days that persists over subsequent weeks, with no evidence of reversals or inter-temporal substitution. This sustained rise suggests a structural shift in behavior, which is consistent with a permanent belief update or change in preferences, rather than salience or coordination.

Second, we investigate spatial diffusion patterns of exposure to large protests, expanding

²We focus on the ever-treated sample, meaning those municipalities with at least one scheduled PEGIDA protest and perform an array of exercises to validate our TWFE approach following Roth et al. (2022). We cluster standard errors at the municipal level.

³To ensure that our findings do not stem from selective policing or reporting of hate crimes we run the following exercises: adding state-by-week and sub-state region-by-week fixed effects allows us to account for changes in policing strategies and funding which are typically mandated at the state level. Leveraging information on police staffing at the state-year level, we find no correlation between cumulative protest size or number of pleasant protest days on police staffing. We also examine whether antisemitic hate crimes, clearance rates or the share of “easily observable” hate crimes increased and find no effect.

the sample to all 10,000 municipalities in Germany. Individuals will update their beliefs about the preferences of others not only based on local protest but also based on support for the movement elsewhere. We consider two channels of diffusion: geographic proximity and social media networks. Specifically, we create a measure of right-wing social media connections based on follow relationships between Twitter users and PEGIDA’s main account, and tracking tweets and retweets containing the word “PEGIDA”. We also account for more general social media connections by creating the equivalent measures based on a random sample of tweets.

This approach allows us to examine the following scenario. Suppose two municipalities, i and j , neither of which hosts a PEGIDA protest in week t , are equidistant from a protest location l and have the same level of general social media exposure to that location. However, municipality i has stronger social media ties to pro-PEGIDA users in l . When l hosts a large protest in t , does hate crime rise more in i than in j ? Our findings indicate that right-wing social media connections to protest-hosting locations significantly amplify local hate crimes, with effects persisting over subsequent weeks. In contrast, broader social media and geographic proximity drive only short-term increases in hate crimes and at less than one-tenth the magnitude. Crucially, neither form of proximity affects local protest occurrence or size, ruling out simultaneity or crowding-out concerns. These results suggest that large protests primarily embolden an already receptive audience in other locations via right-wing social media networks, rather than diffusing support more broadly. It also suggests that the perceived probability of detection and punishment by police, which may drive local hate crimes but not hate crimes in distant locations, explains some of our findings.

Third, we ask whether larger right-wing protests spur counter-mobilization. On the one hand, larger PEGIDA protest should increase issue salience for their opponents and thus drive counter-protest. At the same time, pro-refugee activists may downward correct their belief about local support for refugee issues and thus discourage collective action in opposition to PEGIDA. Using the near-universe of newspaper articles on PEGIDA from the GENIOS database, we identify anti-PEGIDA protests at the municipality level. Replicating our baseline specification reveals no increase in either the likelihood, frequency, or size of counter-demonstrations following larger PEGIDA protests. We then investigate whether this can be explained by heightened costs of counter-mobilization (safety, intimidation or effort), examining the volume of #refugeesWelcome tweets and the net positive sentiment among PEGIDA-related tweets. Our results reveal an uptick in pro-refugee tweets and decline in net positive sentiment among PEGIDA-related tweets. Thus, although protest size raises anti-refugee violence, it does not appear to spur a strong in-person counter-reaction but some pro-refugee mobilization online.

Fourth, we consider the media ecosystem surrounding PEGIDA protests. While coordination may happen on the ground, the signal about other peoples’ preferences can be magnified or muted depending on the coverage of these protests in newspapers and on social media. Constructing a novel database of local newspaper readership and Twitter usage by municipality, we show that local hate crimes rise more sharply when municipalities are exposed to pro-PEGIDA articles and tweets, even after accounting for the direct effect of local protest size. This effect persists into subsequent weeks. Anti-PEGIDA coverage exerts a mildly dampening influence on hate crimes, but it is smaller in magnitude and dissipates more quickly. This asymme-

try suggests that extensive pro-movement coverage amplifies the perceived social acceptance of anti-refugee violence more strongly than any countervailing narratives, underscoring how both traditional and digital media can sustain the emboldening signal of large protests.

Fifth, we employ a large language model (LLM) to classify perpetrator attributes and crime characteristics across nearly 9,000 hate crime descriptions. We find that larger right-wing protests spur a rise in single-offender incidents, consistent with the idea that declining social stigma emboldens individuals to act without group support. This effect is particularly pronounced among perpetrators with documented extremist or recidivist backgrounds, indicating that those already inclined to violence interpret bigger protests as a license to escalate violence. In contrast, there is no increase in group-based offenses or impulsive attacks, suggesting that these crimes reflect enduring changes in perceived norms rather than spur-of-the-moment aggression. Reinforcing this interpretation, the number of attacks that take place in busy, highly visible public venues increases, pointing to lower perceived social costs for offenders. Finally, we detect no increase in explicitly protest-related offenses, implying that the uptick in hate crimes stems from radicals feeling individually encouraged to commit violence rather than any direct spillover from protest activities.

Finally, we consider whether opportunities for coordination are a binding constraint in committing hate crimes, using football matches as a placebo. Football events, particularly those involving far-right fan bases, generate sizable crowds that are plausibly suited for the same face-to-face organization one might expect at political rallies. Thus, if large gatherings in general (rather than an explicit protest message) fuel hate-crime coordination, we would also observe a measurable uptick in xenophobic violence following such matches. Drawing on data for three national leagues plus the domestic cup between 2014 and 2020, however, we find no discernible effect of match occurrence, size, or intensity on hate crimes, even for clubs linked to extremist fans.

Overall, our findings suggest that local protest attendance increases political violence by providing a signal of others' anti-minority sentiment. While marginal attendance may be driven by a desire of moderates to signal support for restrictive migration policy to political elites and policy makers, it also serves as a signal to the political extreme. Social media networks and newspaper coverage amplify this signal, eroding norms against violence without significant counter-mobilization. Consequently, even small initial differences in support for the movement can have long-lasting consequences for public safety due to an unraveling of norms, either through informational cascades (Bikhchandani et al., 1992), "sparks and prairie fires" (Kuran, 1989) or tipping points in peer sanctions (Bernheim, 1994; Kandel & Lazear, 1992).

We situate our study within three key strands of the literature. First, we contribute to a literature showing how minor shifts in perceived social acceptance can trigger cascading changes in collective behavior. Signals from populist leaders can normalize discriminatory views (Bursztyrn et al., 2020; Grosjean et al., 2022; Müller & Schwarz, 2021, 2023; Ajzenman et al., 2023), while social image concerns shape whether individuals publicly endorse or punish such behavior (Andreoni & Bernheim, 2009; Perez-Truglia & Cruces, 2017; Bernhardt et al., 2018). In addition, theoretical models illustrate how small initial differences can yield sharply divergent outcomes, be it through herding on early movers (Bikhchandani et al., 1992), "sparks and prairie fires"

(Kuran, 1989; Correa et al., 2025), or tipping points in peer sanctions (Kandel & Lazear, 1992; Bernheim, 1994; Boyer et al., 2024). Our findings build on these insights by showing that protest attendance, amplified through social and traditional media, can lower the perceived social cost of extremist acts.

Second, we add to prior work which shows that protests can shape mainstream political outcomes, often boosting electoral support for parties aligned with the protesters' agenda (Madestam et al., 2013; Larrebourg & González, 2021; Casanueva, 2021; Lagios et al., 2025; Cantoni et al., 2024). These studies highlight the influence of social movements on democratic electoral shifts; however, their focus on votes and parties necessarily overlooks effects on the political extremes. In particular, electoral outcomes are bounded within democratic norms and do not capture radicalization beyond those limits. Building on this literature, our study examines protest effects at the extreme right of the political spectrum, analyzing how protests might fuel radicalization outside the usual electoral realm.

Finally, we contribute to the literature on hate crimes and their determinants. Surges in hate crimes have been linked to periods of economic distress and insecurity (Bray et al., 2022; Jaschke et al., 2022), as well as to moments that shift societal norms of acceptability (for instance, triggering events that embolden prejudice) (Hanes & Machin, 2014; Romarri, 2020; Carr et al., 2020). In some cases, economic and normative factors interact – for example, sudden demographic changes or refugee inflows can heighten both economic anxiety and social tension, leading to spikes in hate incidents (Entorf & Lange, 2019; Han et al., 2023; Dipoppa et al., 2023). Recent work also highlights the role of online networks in catalyzing hate crimes, as extremist content on social media can translate into real-world violence (Müller & Schwarz, 2021, 2023; Levy & Mattsson, 2023). Extending these insights, we show that social media networks and newspaper coverage of offline right-wing protests act as a powerful precursor to hate crimes.

2 Background and Data

2.1 Background

Refugee influx to Germany and hate crimes. Germany has emerged as a primary destination for refugees in Europe, with over 1.6 million asylum applications filed between 2015 and 2018 alone, representing more than 40% of all applications in the European Union during this period (Eurostat, 2019). The surge in asylum applications can be attributed to the eruption of the civil war in Syria and the growing threat of the so-called Islamic State in Iraq, as well as political and social unrest in other parts of the Middle-East and Sub-Saharan Africa leading to a movement of hundreds of thousands of refugees from Syria, Iraq, Afghanistan as well as from Albania, Kosovo and Eritrea.

The peak of asylum applications in Germany occurred in late 2015, following Angela Merkel's controversial decision to admit refugees stranded in Hungary. This decision deviated from the Dublin Regulation, which assigns responsibility for administering an asylum request to the country of first-entry. However, the regulation was effectively abandoned before September 2015, as registration and administrative capacities in Italy and Greece were overwhelmed by immigration pressures, and refugees sought to move to Northern Europe. In March 2016, the

European Union established a treaty with Turkey to encourage stricter controls by Turkish authorities at its Western shores, leading to a decline in asylum applications in Germany, which have remained relatively low since then.

In the early stages of the refugee crisis, Germany showed a strong sense of *Willkommenskultur* or “culture of welcome,” with many Germans volunteering to help refugees and participating in demonstrations in support of their cause. However, as the number of refugees increased, this sentiment began to shift. Some Germans expressed concerns about the economic and social impact of refugees, with right-wing parties and anti-immigrant groups gaining momentum. The issue became highly politicized, with debates surrounding the government’s handling of the crisis and calls for tighter immigration policies. Alongside the success of PEGIDA and the right-wing party AfD, violence against minorities and particularly refugees began to rise and peaked in 2016.

Right-wing protest under PEGIDA. The Patriotic Europeans Against the Islamisation of the Occident, or PEGIDA, movement was founded by Lutz Bachmann in late 2014. It originated in Dresden, the capital of the state of Saxony in Eastern Germany, as a local Facebook initiative with approximately 300 participants in the first demonstration. The movement grew exponentially, following the influx of refugees to Germany in 2015 and reached its peak in mid-January 2015 with approximately 25,000 participants in Dresden. The success was accompanied by offshoots in other cities within and beyond Germany (Berntzen & Weisskircher, 2016).

The movement referenced the renowned Monday demonstrations that took place in the former German Democratic Republic (GDR) in 1989. These demonstrations have since become a symbol of peaceful civic engagement and political change in the minds of many Germans. PEGIDA has appropriated the concept of these demonstrations in an effort to portray itself as a concerned citizens’ movement calling for significant reform in immigration policy, as well as the protection of the Christian-Jewish tradition in Europe. Since its inception, PEGIDA has adhered to a consistent three-part structure every Monday, starting with a round of speeches, followed by an evening stroll, and concluding with a closing rally.

Over time, PEGIDA sharpened its profile as a nationalistic, xenophobic and Islamophobic movement with ties to neo-Nazis and other fascist groups (Vorländer et al., 2018). The movement aimed to fuel xenophobic sentiment in the population and gained electoral influence with the rise of a new right-wing populist party - the Alternative for Germany (AfD) in 2016. In 2021, the German domestic intelligence service (*Verfassungsschutz*) has classified the goals of the movement as unconstitutional. Its founder, Lutz Bachmann, was sentenced to two years of probation in 2020 for inciting hate at PEGIDA protests.

The case of Heidenau Heidenau, with a population of 16,000, found itself thrust into the spotlight on August 21, 2015, when local authorities, responding to the escalating refugee crisis in Europe, converted an empty hardware store into a temporary home for asylum seekers. This decision was met with immediate and organized opposition, predominantly fueled by the National Democratic Party of Germany (NPD), known for its right-wing extremist views and for their close ties to PEGIDA. Heidenau is only a 20 minute car ride away from Dresden, PEGIDA’s stronghold in Eastern Germany.

At the peak of PEGIDA’s success in late summer of 2015, a group of Neo-Nazis rallied around and attacked the refugee camp. The surge in violence led to clashes with the police and sparked a national political debate. The involvement of federal authorities, including the visit by vice chancellor Gabriel and cautionary remarks by chancellor Merkel, meant to emphasize the government’s commitment to defending refugee rights and combating extremism. However, these high-profile interventions also sparked controversy, with some critics arguing that the government’s response was either too late or insufficiently forceful to deter future xenophobic incidents.

Media outlets rushed into the city to document what was considered ”a look into the psyche of the country”, quoting – for instance – a witness about seeing ”faces that are known here in Heidenau. When the 200 turned up, there were people standing on the railway embankment, who were cheering and clapping. Elderly with bicycles, children were there too. They clapped as if at a summer movie night, as the right-wingers moved towards the Praktiker [hardware store]. And in the sports store, the baseball bats were sold out”.⁴

2.2 Main Data Sources

Protest. We take data on PEGIDA protests from the Right-Wing Extremist Mobilization in Germany data set, created by Kanol & Knoesel (2021). The dataset on right-wing demonstrations in Germany between 2005 and June 2020 was created using information from the German federal government’s responses to parliamentary questions tabled by the opposition left-wing party *Die Linke*. The dataset includes information on the location, date, number of participants, organizing actors, and motto of the right-wing extremist demonstrations.

For our purposes, we restrict attention to PEGIDA Monday protests specifically because of their ritualized and pre-planned nature and because these were explicitly anti-refugee. These represent the vast majority of PEGIDA protests (see Figure A.1). We focus only on Monday protests between 2015 and the end of 2019, leaving us with a total of 398 protest days with an average of 150 and a maximum of 2,300 participants (see Table A.1). Many of the protests were located in the former Eastern part of Germany but we show in Figure 1 that we find offshoots in Western Germany, including the biggest states North-Rhine-Westphalia and Bavaria in the Southern and Western part of the country.

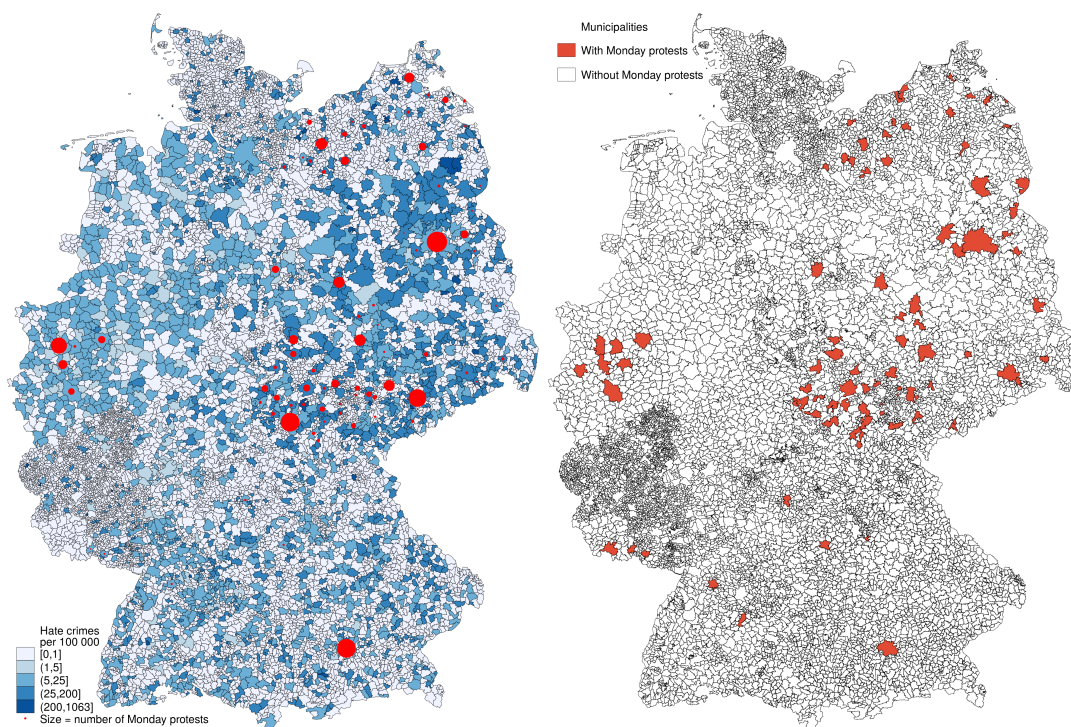
Hate Crimes. We scrape data on hate crimes from the chronicle reported by the Amadeu Anotonio Foundation (AAF) and PRO ASYL Foundation, for the period of 2015-2020. Their data is taken from various sources, including newspaper articles, police press releases, reports from local crime registries and community centers for those affected by right-wing, racist and anti-Semitic violence. Similar to the data on right-wing protests, the most common source are governmental answers to inquiries made by the Left party. Since 2014, every quarter, the Left party in Germany poses a brief parliamentary question (*Kleine Anfrage*), asking the Federal Government to list all cases of attacks directed at refugees or their accommodation, which are considered by the police as right-wing politically motivated crimes (*PMK*). For each case, the government reports its date, location, and the type of crime committed.

⁴From one of the main news outlets Die Welt on September 1st 2015.

We distinguish between antisemitic hate crimes and hate crimes committed against “visible minorities” for two reasons. First, PEGIDA claims to defend the Christian-Judaeo tradition in Germany against Islamization, targeting immigrants from Muslim-majority countries specifically. However, political analysts argue PEGIDA’s ostensibly pro-Jewish stance acts as a fig leaf to legitimize a xenophobic agenda that remains fundamentally intolerant (Vorländer et al., 2018). Second, antisemitic hate crimes fall enhanced punishment in Germany with specific protections granted.⁵ It is worth noting that in the context of hate crimes against other minorities, the vast majority of recorded hate crimes in our data set are related to physical assault or arson (typically of refugee camps), rather than hate speech. We provide examples of hate crimes in Appendix C and describe in detail how we use Large Language Models to parse information on the type of hate crime from 9,000 hate crime descriptions.

We assign hate crimes to each municipality and week, distinguishing between the weekdays on which these hate crimes were committed. Overall, there are approximately 1,800 week-municipality observations with at least one hate-crime committed against refugees. We show in Figure 1 that the cumulative number of hate-crimes per 100K inhabitants across municipalities between 2015 and 2020 is spread evenly across the country. We show in Figure 2 that the number of hate-crimes per week and the number of protests per week follow similar patterns, with hate-crimes lagging behind by about one to three weeks.


Figure 1. **Hate crimes and PEGIDA protest across municipalities**



Note: Left panel shows cumulative number of hate crimes (blue shades) and PEGIDA protests (red circles) across municipalities over entire observation period (2015-2020). Right panel shows sample of ever-treated municipalities in red, that is municipalities with at least one Monday PEGIDA protest between 2015 and 2020.

⁵Section 130 of the Criminal Code (Volksverhetzung or incitement of the people) criminalizes incitement to hatred against segments of the population and explicitly bans Holocaust denial and trivialization.

Figure 2. **Hate crimes and PEGIDA protest over time**



Note: Figure shows two-week moving average of the number of scheduled Monday PEGIDA protests and the number of hate crimes across all municipalities.

Social Media. We use four measures to proxy social media use at the local level: *i*) overall Twitter use, *ii*) PEGIDA tweets, *iii*) followers of the official AfD Facebook account and *iv*) pro-refugee tweets. First, we develop a measure for Twitter usage for each NUTS-3 region in 2013 and 2014 based on a random sample of 600,000 tweets. We geolocate authors using the location indicated in their profile.⁶ In addition, we collect all tweets in German and in English containing the word PEGIDA posted between October 2014 and 2021. This dataset consists of 2,068,258 (and 659,709 geo-localized) tweets and retweets, along with their date of posting, their retweet status, the text of the tweet, and information about the author. We also collect the list of followers of the main PEGIDA Twitter account, and the list of followers of this account. Based on this information we build three indices of social media proximity, which we describe in more detail in section 4.2. Information on the number of followers of the Facebook page of AfD prior to 2015 are taken from Müller & Schwarz (2021).⁷ Lastly, we proxy pro-refugee sentiment, using all tweets and retweets in German mentioning the hashtag #RefugeesWelcome between 2013 and 2018. We are able to geo-localize 150,000 of about 390,000 tweets.

⁶We use the Twitter Academic Search API to collect all Twitter data. We pick 6,000 random instants during this period and collect 100 tweets and retweets in German at each instant. Since the Twitter API does not allow to search directly for all tweets in German, we search for tweets containing the 100 most frequent words in German, as listed by Sharoff (2006) on the website <http://corpus.leeds.ac.uk/frqc/>. The Twitter API gives users' location at the time tweets were collected, not posted). We use the Nominatim geocoder from the OpenStreetMaps project to associate the location field to geographical coordinate, and remove locations outside of Germany, as well as locations that are too general (e.g. "Germany" or "Bavaria"). This gives us an estimate of the rate of tweets posted at each instant from each region (expressed as tweets per second), which is then aggregated at the region-year level.

⁷This data is localized at the collective municipality (*Gemeindeverband*) level. These groups may include multiple municipalities. We map the (per capita) number of followers to all municipalities within a group.

Newspapers. We build a set of articles mentioning pro-PEGIDA protests and counterprotests. We use this dataset as an alternative source of pro-PEGIDA protest information, to construct a measure of offline counter-protest, and to measure the exposure to articles talking about PEGIDA protests and counter-protests. We use the GENIOS newspaper database and search for all articles mentioning PEGIDA in the 2015-2019 period. We use a large-language model to analyze the selected articles and extract information about protests and counter-protests mentioned there, then aggregate the information from all these articles to build the full dataset. The construction is detailed in Appendix section C.6.

Additional data sources. Regional controls were retrieved from four administrative sources. Labor market data were taken from the Federal Employment Agency, election outcomes from the Federal Returning Officer (Bundeswahlleiter), and the rest from the Statistical Offices of the Federal States (Statistische Ämter des Bundes und der Länder) and Federal Criminal Police Office (Bundeskriminalamt). Appendix Table C.1 describes all regional controls, their geographic granularity, time coverage and frequency, as well as the data sources. AfD vote share, population density, age structure of population, share of females, share of foreigners unemployed, and share of unemployed are available at the municipality level. Refugee share, share of asylum recipients, share of foreigners with academic qualification, and GDP per population are available at the district-level (*Kreise*). We also use yearly police staffing data at the state level from We use NUTS3 boundaries of 2013 to harmonize administrative changes over time, which we describe in more detail in Appendix C.

2.3 Descriptive Statistics

Table A.1 reports summary statistics for our sample of interest, i.e. municipalities with at least one PEGIDA Monday protest and averaged over all years from 2015 to 2020. Roughly 8% of municipality-week observations feature at least one hate crime against refugees, while about 2% have a Monday PEGIDA protest. Conditional on a protest taking place, the mean number of participants is around 147, but can reach as high as 2,300. Economic and demographic measures also highlight substantial diversity. The mean GDP per capita at the municipality level is just above 31,000 euros, while the richest municipalities exhibit per capita GDP levels of more than 2.5 times the average. Population density averages around 715 inhabitants per square kilometer, reaching up to roughly 4,736 in urban areas. On average, 4.4% of the labor force is unemployed, and about 1.4% of the population are registered refugees, but these shares vary considerably across regions and can reach up to 5%. Regarding political and social indicators, the right-wing vote share for the AfD stands at around 14% in these municipalities, and the share of total crime cases over 100,000 inhabitants averages 0.65%, again with notable heterogeneity. We also report descriptive statistics for the sample of municipalities with no PEGIDA protest. The probability of recording a hate-crime in ever-treated municipalities is almost 40 times higher than in other municipalities. As expected, the ever-treated municipalities are more densely populated, have a higher unemployment rate, a higher vote-share for the right-wing party but a lower overall crime rate and host a similar number of refugees relative to their population size.

3 Research Design and Main Results

3.1 Estimating Equation

We start by investigating the relationship between right-wing protest participation and hate crimes, using a two-way fixed effects approach that covers the period between 2015 and 2020. Specifically, we exploit municipality (N=10K) and week (T=251) variation to estimate a linear probability model and of the following form:

$$HC_{it} = \beta P_{it} + \gamma X_{it} + \delta_t + \mu_{im} + \nu_i \times t + \epsilon_{it} \quad (1)$$

Our outcome of interest HC_{it} is a binary variable that takes the value 1 if any hate crime was recorded in municipality i in week t and 0 otherwise. Our coefficient of interest β captures the effect of protest size P_{it} , measured as the log of one plus the number of participants at a scheduled Monday PEGIDA demonstration, on the probability of observing at least one hate crime in the same week. We include week fixed effects δ_t , municipality by month of the year fixed effects μ_{im} , as well as municipality specific linear time trends $\nu_i \times T$. X_{it} denotes a large battery of time-varying municipality-level controls, which we summarize in Table C.1 and describe in more detail below. Standard errors are clustered at the municipality level. Throughout, we focus on the ever-treated sample, i.e. municipalities that experienced at least one protest during the observation period.⁸

This research design has several advantages. First, we can account for time varying factors that are common to all municipalities, such as the overall popularity of the right-wing movement, national and European election cycles or the overall salience of the refugee issue. Second, the model absorbs any unobserved heterogeneity at the municipality level, including - for instance - the root determinants of anti-immigrant sentiment. The municipality by month of year fixed effects capture any seasonal differences intrinsic to municipalities that could be related to protest size and the likelihood of committing hate crimes. For instance, some municipality may host events during certain months of the year that appeal to a right-wing crowd which also participates in protests and commits hate crimes. Municipality-specific linear time trends capture gradual changes within municipalities over time, such as shifting demographic compositions, economic conditions, or long-term changes in social attitudes.

Our specification also accounts for state-dependence in protest participation and hate crimes. It is possible that a large protest may encourage subsequent protest participation. Similarly, hate crimes may normalize future violence against minorities. Another concern relates to inter-temporal substitution and anticipation effects, where individuals that intend to commit hate crimes strategically choose the timing.⁹ In all cases, β would be upward-biased through the persistence, serial correlation or inter-temporal substitution of violence and protest size. Therefore, the set of controls X_{it} includes the likelihood of observing a hate crime or a protest in the

⁸Recent developments in the literature emphasize caveats in the classical difference in differences setting when it involves both group and time fixed effects, i.e. Two-Way Fixed-Effects or in short TWFE (e.g., De Chaisemartin & D'Haultfœuille (2022), Wooldridge (2021), Roth et al. (2022), Goodman-Bacon (2021)). We address the issue of forbidden comparisons in the presence of staggered treatment more carefully in Appendix ???. Let us preview here that the TWFE estimation does not produce any negative weights and is robust to using the difference in differences estimate of ?.

⁹We show in an event-study setting that anticipation or inter-temporal substitution do not play a role.

previous period as well as the lagged number of protest participants.

To account for potential confounding factors and gain precision in our estimates, we include several municipality level time varying controls that proxy the overall economic and political conditions in the municipality as well as the propensity to commit crimes. We account for socio-economic conditions by controlling for GDP per capita, population density and the share of unemployed in the municipality. Hate crimes against minorities require both the presence of minorities and the presence of xenophobic individuals. Hence we include the share of refugees in the municipality as well as the vote share for the right-wing party AfD in the latest national or European election. We also control for the overall propensity to commit and record crimes in the municipality with the total number of documented crimes per 100K inhabitants.

3.2 Instrumental Variable Strategy

Weather on scheduled protest days

In the previous section, we have outlined how the set of controls and fixed effects accounts for important sources of unobserved heterogeneity. However, it is still possible that protest size is correlated with unobserved factors that also influence violence against minorities. For instance, those that intend to commit hate crimes may also be involved in the organization of the protest and therefore the mobilization of protesters. In the case that these variables change within municipalities over time, β would not capture the causal effect of protest size but the far right mobilization potential of a municipality in the form of encouraging protest participation and hate crimes.¹⁰

In order to address this concern, we employ an identification strategy that relies on exogenous variation in local weather conditions at a given protest day. Figure 3 depicts (residualized) protest participation in relation to rain and temperature in the municipality on the protest day: participation follows an inverse U-shape, indicating that more protesters take to the street in moderate temperatures and that protest participation decreases with higher levels of precipitation. We define a variable that captures pleasant weather, indentifying the appropriate weather conditions for protest participation in Appendix Figure A.2. We report the coefficients for separate regressions that estimate the effect of different rain and temperature cut-offs interacted with the protest dummy on the log number of participants to show that rain above 10 mm and temperatures above 21 degrees Celsius are associated with significant drops in the number of protest participants.¹¹ We combine these two components into a dummy variable for pleasant weather during protest times (between 12 pm and 5 pm) that switches on if temperatures are between 0 and 21 degrees Celsius and there is no heavy shower (average precipitation of less than 10 mm per square meter). Approximately 25% of all protest happen on days with unpleasant weather. In a series of exercises presented in section 3.4, we show that results are not sensitive to changes in the cut-off value and that they hold using LASSO-selected instruments. Our first stage takes the following form:

¹⁰Note that if the right-wing mobilization potential of municipalities changed over time in a linear fashion the OLS specification would account for this with linear time trends.


¹¹In these regressions we include again municipality-month of the year fixed effects as well as week fixed effects, control for protest and weather cut-off separately and include the interaction between the full set of controls and protest and weather cut offs.

$$\begin{aligned}
P_{it} = & \alpha \text{ protest}_{it} \times \text{weather}_{it} \\
& + \eta_1 \text{ protest}_{it} \times X'_{it} + \eta_2 \text{ weather}_{it} \times X'_{it} \\
& + \theta_1 \text{ protest}_{it} + \theta_2 \text{ weather}_{it} + \theta_3 X_{it} \\
& + \nu_i \times T + \mu_{im} + \delta_t + \epsilon_{it}
\end{aligned} \tag{2}$$

We estimate the differential effect of a protest during pleasant weather (excluded instrument) on protest participation. This approach also allows us to control for protest events and weather conditions separately. In addition, we condition on the full set of control variables X'_{it} and their interaction with the pleasant weather dummy ($\text{weather}_{it} \times X'_{it}$) and protest ($\text{protest}_{it} \times X'_{it}$). Equivalent to our two-way fixed effects estimation, we include week fixed effects, municipality-specific linear time trends, as well as municipality month of the year fixed effects to account for seasonal weather differences across municipalities, thereby exploiting deviations from average temperatures and precipitation.


It is worth noting that our instrument variable approach is especially well-suited for examining a signaling mechanism: local weather conditions primarily affect those on the margin of participation, thereby shifting the perceived size and success of a protest in a way that is plausibly unrelated to other drivers of hate crimes. By capturing exogenous variation in the scale of public demonstrations (rather than in whether protests occur at all) we are thus better able to isolate the role of the protest's public signal. Of course, this presumes that there is uncertainty about the elasticity of protest attendance with respect to weather, i.e. individuals do not fully internalize the weather-driven portion of attendance, interpreting higher turnout as genuine backing rather than incidental.

Figure 3. **Residualized protest participation and weather conditions**



Note: Left panel shows the binned scatterplot and fitted line with 95% confidence intervals between the log of 1 + the number of participants at a Monday PEGIDA protest net of baseline controls and fixed effects on the (binned) precipitation measured as an average rain in mm per sqm between 2pm and 5pm on the protest day. The right panel repeats this analysis, this time using the average temperature in Celsius between 2pm and 5pm on the protest day.

Figure 4. **Time variation of PEGIDA protests on pleasant days**



Plausibility of exclusion restriction and first stage

A causal interpretation of the coefficient requires that the interaction between weather and protest only impacts hate crimes through its effect on protest participation.¹² Exploiting exogenous variation in the interaction between protest events and pleasant weather allows us to address several endogeneity concerns. First, determinants of hate-crimes that are related to any right-wing protest itself will be captured in θ_1 . For instance, the most extreme fraction of the right-wing movement that participates in every scheduled protest (irrespective of weather conditions) could use the occasion to conspire and coordinate hate crimes against minorities. Second, any weather conditions that are conducive to crimes more generally and hate-crimes more specifically, will be captured in θ_2 (Heilmann et al., 2021; Field, 1992). Controlling for weather conditions separately alleviates many of the concerns associated with a violation of the exclusion restriction in the context of weather data (Sarsons, 2015; Mellon, 2021).¹³ Moreover, heterogeneous responses to protest and pleasant weather that relate to municipality characteristics are captured in η_1 and η_2 , respectively. Compared to similar instruments in the literature, which typically use weather conditions to instrument for the occurrence of protests rather than their scale (Madestam et al., 2013; Beraja et al., 2023; Qin et al., 2024), our approach benefits


¹²Participation may be accompanied by increased public attention for right-wing protests. For instance, protests on pleasant days may attract relatively more journalists or generate more positive imagery of the protest. This does not invalidate our empirical strategy since we are interested in the emboldening effect of seemingly successful right-wing protest. This can be in the form of protest participation, in related positive protest coverage (Zhong & Zhou, 2012) or in an increase in the positive mood of people when participating in or learning about the protest (Goetzmann et al., 2015; Jiang et al., 2022). We assess this possibility in more detail in section 4.3.

¹³We also show later that controlling for weather on the day of the crime does not change our results, assuaging concerns about the serial correlation in pleasant weather.

from a richer set of controls and fixed effects. This framework not only strengthens the plausibility of the exclusion restriction but also mitigates endogeneity concerns by isolating the effect of protest size from confounding factors associated with weather and protest.

Figure 5 presents a plausibility check for the assumption that protest during pleasant weather is exogenous to unobserved municipality characteristics. Specifically, we predict the probability of treatment with an array of municipality-level characteristics. These characteristics are standardized and categorized into demographics (age and gender distributions), socioeconomic status (employment, education, and income metrics), political preferences (party vote shares and voter turnout), cultural/religious composition (shares of religious affiliations), migration-related factors (asylum, refugee, and migration shares), and crime statistics.¹⁴ The coefficients demonstrate no systematic correlation with the treatment, reinforcing the plausibility of our identifying assumption. Importantly, the fixed effects and controls of our baseline specification capture all characteristics presented in Figure 5 and others that could be related to selection into treatment.

Figure 5. Municipality characteristics and protest on pleasant days



Note: Plausibility of the instrument exogeneity. We estimate the following regression: $\sum Pleasant \times Protest_{muni.year} = \beta X_{muni.year} + \sum Pleasant_{muni.year} + \sum Protest_{muni.year} + \gamma_{state} + \theta_{year} + \epsilon_{muni.year}$. We regress the set of baseline control variables and additional municipality characteristics on the cumulative number of protests on pleasant days in a municipality and year. We include state fixed effects as well as year fixed effects and control for the cumulative number of pleasant days and cumulative number of Monday protests separately. Standard errors are clustered at the state level, and 95% confidence intervals are represented. Coefficients are standardized for comparability. We re-scale population density and GDP per capita in a 1:10 ratio for readability.

We present the first stage results in Table A.2, reporting the coefficients for the interaction between pleasant weather and protest, as well as protest and weather separately. Columns 1 to 4, successively add the set of controls X_{it} , including their interaction with weather and

¹⁴The model, described by $\sum Pleasant \times Protest_{muni.year} = \beta X_{muni.year} + \sum Pleasant_{muni.year} + \sum Protest_{muni.year} + \gamma_{state} + \theta_{year} + \epsilon_{muni.year}$, regresses a comprehensive set of municipal characteristics against the count of protests on pleasant days per municipality and year. State and year fixed effects are incorporated, along with controls for the number of pleasant days and Monday protests. Standard errors are clustered at the state level.

protest. Throughout, we find that the interaction term significantly and positively predicts protest participation. The magnitudes remain statistically indistinguishable across columns while the first stage becomes weaker (but remains above the conventional threshold) when including additional controls. Protest occurrence alone is, expectedly, the strongest predictor for participation while weather alone does not predict protest size. The coefficient of the interaction in column 4 suggests that a protest during pleasant weather increases the number of participants by 25 percentage points.

3.3 Protest size and hate crimes

Table 1 presents the main results, reporting OLS estimates in Panel A, 2SLS estimates in Panel B, and reduced form estimates in Panel C. Column 1 includes baseline controls, such as GDP per capita, population density, unemployment share, and lagged values for participants, protests, and hate crimes, as well as week fixed effects, municipality-by-month-of-year fixed effects, and municipality-specific linear time trends. Column 2 adds the share of refugees, capturing the supply side of potential victims. Column 3 accounts for municipality-level political characteristics, specifically the vote share for the far-right AfD party in the most recent election to capture the potential supply side of perpetrators. Finally, column 4 incorporates the local crime rate per 100,000 inhabitants to account for police resources and crime reporting. Note that the 2SLS estimates and reduced form estimates also include the interaction between these controls and protest as well as weather separately.

Across all specification, the coefficient of interest remains positive and statistically significant, suggesting that larger protests increase the incidence of hate crimes against minorities. The 2SLS estimate in column 4 suggests that a 1% increase in protest size increases the probability of observing at least one hate crime in the same municipality and week by approximately 0.5 percentage points. To better contextualize this effect, consider a 25% increase in protest size, which is the average effect of a pleasant day on protest participation: the likelihood of a hate crime would increase by approximately 11 percentage points. This is a substantial change relative to the baseline probability of 8.1%.

Throughout, the 2SLS estimates exhibit larger coefficients than the OLS estimation. One possible explanation relates to the compliers in this context. Specifically, marginal participants, i.e. those whose attendance is influenced by exogenous weather variation, may be critical in generating a surge in hate crimes. It is possible that endogenous changes in participation over time are already accounted for in the expectations of radicals. Exogenous variation in protest participation through weather may act as a surprise, amplifying the public signaling effect of the protest and further emboldening individuals to commit hate crimes. This may be exacerbated by the composition of these marginal participants who - presumably - do not belong to the most radicalized fraction of society. If salience or coordination costs were the primary forces, then we would not expect the same pattern.¹⁵

¹⁵Let us caveat, however, that the difference in OLS and 2SLS estimates could also be driven by other factors, including differences in LATE and ATE or measurement error in the treatment variable.

Table 1. **Protest Participation increases probability of hate crimes**

	Any hate crime in the same week			
	(1)	(2)	(3)	(4)
Panel A: OLS				
Log(participants)	0.0196*** (0.00500)	0.0194*** (0.00496)	0.0192*** (0.00493)	0.0192*** (0.00493)
Panel B: 2SLS				
Log(participants)	0.400*** (0.146)	0.504** (0.216)	0.513** (0.218)	0.513** (0.215)
Panel C: Reduced Form				
Pleasant weather × Protest	0.127*** (0.0326)	0.127*** (0.0319)	0.126*** (0.0309)	0.128*** (0.0308)
Protest	0.283** (0.120)	0.285** (0.142)	0.290** (0.138)	0.317** (0.139)
Weather	-0.0243 (0.0209)	-0.0237 (0.0210)	-0.0195 (0.0318)	-0.0193 (0.0328)
Observations	21,678	21,678	21,678	21,678
Municipalities	84	84	84	84
Mean dep. var.	0.0810	0.0810	0.0810	0.0810
F first	20.26	10.76	10.54	10.77
Week FE	Yes	Yes	Yes	Yes
Municipality × MoY FE	Yes	Yes	Yes	Yes
Municipality linear trends	Yes	Yes	Yes	Yes
Lagged controls	Yes	Yes	Yes	Yes
Economic controls	Yes	Yes	Yes	Yes
Refugee share		Yes	Yes	Yes
Right-wing share			Yes	Yes
Crime rate				Yes

Note: Differences in differences estimation with week and municipality-month of year fixed effects and municipality linear time trends. Ever-treated sample only. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Outcome comes from ProAsyl and AAF and is measured as a dummy variable equals to 1 if any hate crime against refugees was committed between Monday and Sunday of the same week. Instrument in Panel B is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleinbergen-Paap F-Statistics are reported at the bottom of the table. Controls in column 1 comprise GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, columns 2 to 5 add refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants respectively, all measured at the end of each calendar year. Panel B and C additionally include the interaction between controls and the weather as well as the protest dummy.

3.4 Robustness Checks

Bias in reporting of hate-crimes

One potential concern with our analysis is differential reporting of hate crimes if, after large protests, the population becomes more inclined to report such incidents or the police become more vigilant in recording them. Either channel could lead to an increase in reported hate crimes that reflects heightened attention rather than an actual increase in offenses.¹⁶ To empirically test for reporting bias, we conduct several exercises in Table A.3. First, because policing strategies and funding are determined at the state level, if state-level leadership systematically changed scrutiny or recording of hate crimes, we would not expect within-state variation in hate-crime incidence after large protests. Including both state-by-week (column 1) and sub-state region-by-week (column 2) fixed effects does not alter the estimates, suggesting no systematic state-wide reporting changes.

Second, we examine whether other hate crimes against other minorities increased in the aftermath of large protest, looking at antisemitic hate crimes from the same data source (*Amadeu Antonio Foundation*). It is worth noting here that PEGIDA purportedly “defends the Judeo-Christian tradition” of Germany, emphasizing that the movement is not antisemitic. Larger PEGIDA protests do not increase anti-semitic hate crimes. Next, we check whether a more localized police or community response in municipalities with large protests could inflate recorded hate crimes. If so, we would expect higher clearance rates (i.e., the ratio of hate crimes with identified suspects to total hate crimes) because of more active policing. Columns (3) and (4) show no significant effect on the likelihood of a crime having a suspect or on the clearance rate, respectively. As described in section 2, the threshold for a hate crime classification is comparatively high in Germany (except for antisemitic hate crimes, which we will examine in the next section) and thus less likely to suffer from reporting bias at the margin. Nonetheless, we examine the share of plausibly detectable hate crimes (column 5). Arson is typically more severe, almost exclusively affects refugee camps and is thus easily observable and consistently reported. If the share of these hate crimes over all reported hate crimes decreases, this could be indicative of shift towards more sensitive crime reporting. Reassuringly, we find no evidence for a shift away from more easily observable hate crimes.

Finally, we draw on state-level staffing data for 2015 to 2020. Columns (6) and (7) regress the number of police employees (in 1,000 full-time equivalent positions) on protest size and the number of protests on pleasant days, respectively, while controlling for state and year fixed effects. Again, we find no significant relationship, suggesting that changes in police resources are unlikely to drive our results. Let us also preview here that our evidence on spatial diffusion patterns of hate crimes in the next section speaks against reporting bias.

¹⁶Notably, gradual investments in police resources or increased awareness among citizens would be picked up by municipality-specific linear time trends; any systematic or municipality-specific seasonal differences in crime reporting are captured by the municipality-by-month-of-year fixed effects; and heightened police response to any protest, or differential police response in municipalities with higher levels of recorded crime (including both overall crimes and hate crimes in the previous week) after any protests would be captured by interaction terms with the protest dummy.

Alternative instruments

Next, we examine the possibility that our results are sensitive to the definition of pleasant weather. Figure A.3 reports 2SLS estimates alongside the first stage Kleinbergen-Paap F-statistics for alternative weather cut-offs. We reduce the rain cut-offs by 1 to 3 mm rain per square meter, and reduce the maximum and minimum temperature cutoff between 1 to 3 degrees Celcius. In addition, we use continuous rain and temperature to predict protest size. The F-statistics for the majority of instruments remain around the conventional thresholds of 10 and the estimated second stage coefficients are positive and significant, and of similar magnitude to our baseline result. Continuous rain and temperature deliver weak first stages, potentially because participation is not linearly decreasing along those margins (see Figure 3).

To further decrease the researchers' degree of freedom as well as the risk of finding false positives, we implement a LASSO IV approach similar to Beraja et al. (2023). We take hourly information on precipitation and temperature between 8 am and 9 pm on the protest day and allow each variable to interact with each other and with the protest dummy. We implement a LASSO regression that selects predictors of protest participation based on over 1200 variables. Standard errors are calculated using the cross-fit partialing-out LASSO IV algorithm following Chernozhukov et al. (2018). Figure A.4 reports the second stage coefficients for the LASSO selected instruments. In the top bar, we do not impose that LASSO selects any fixed effects or controls, the second bar imposes the fixed effects structure of our baseline regression and the third bar imposes the inclusion of our baseline controls. Throughout we find a positive and significant association between weather predicted protest size and hate crimes.

Summary of additional robustness checks

We run an array of additional robustness checks, which we summarize briefly here and in more detail in Appendix B. First, we address concerns about spatial correlation and spillovers. Columns (2) and (3) of Table B.1 repeat the analysis at the NUTS-3 and NUTS-2 levels, clustering standard errors at these more aggregated units, while column (1) reports the baseline specification for reference. Columns (4)–(7) then employ Conley standard errors with distance thresholds of 25 km, 50 km, 100 km, and 150 km to account for potential spatial spillovers. Despite minor changes in precision, the magnitudes and significance of the estimated coefficients remain broadly stable.

Second, we test the robustness of our results to sample composition and additional controls. In Table B.2, we expand the baseline controls to include demographic characteristics (e.g., shares of unskilled individuals, women, and age groups), indicators of immigrant and refugee vulnerability (e.g., shares of tertiary-educated, unemployed, and asylum-granted foreigners), and social media measures (e.g., 2014 Twitter penetration, #refugeesWelcome tweets per capita, and Facebook AfD followers before 2015). In Table B.3, we examine the robustness of our results with respect to different fixed effect structures. In Figure B.1, we plot the distribution of treatment effects when dropping single municipalities and single weeks. Throughout, our results remain robust to these changes.

Third, recent work in the two-way fixed effects literature emphasizes that multiple or staggered treatments can lead to biases, including negative weights (Roth et al., 2022). To address

concerns about negative weights in two-way fixed effects (TWFE) estimators (De Chaisemartin & D’Haultfoeuille, 2022; ?), we employ their diagnostic procedure to identify whether the estimated average treatment effects (ATEs) might be biased by negative weighting. Table B.4 reports the number and sum of positive versus negative weights across progressively richer sets of controls. In all specifications, we observe only a small number of negative weights, and their total contribution (sum of weights) is negligible relative to the positive weights. In addition, restricting the analysis to the very first protest event avoids these complications by ensuring only one treatment time per unit. We verify that our results hold using the first PEGIDA protest on a pleasant day in an event-study in Figure B.2. This also allows us to verify that municipalities that radicalize more quickly (i.e., are on different hate crime trajectories) are not more likely to mount their first PEGIDA protest early on. We investigate the timing of hate crimes and dynamic treatment effects in a series of event-studies in more detail in the next section. Figure B.3 shows results when controlling for weather on the day of the hate crime and when controlling for the cumulative number of past PEGIDA protests.

Lastly, we probe the robustness with respect to the definition of the outcome and treatment variables. First, Table B.5 reports coefficients for protest size when excluding hate crimes that were committed on the day of the protest. Next, Table B.6 examines the sensitivity of our results with respect to potential non-linearities in the treatment variable. Instead of the log-transformed number of participants, we take the inverse hyperbolic sine transformation, the number of participants as a share of the overall population, as well as the absolute number of participants. In addition, instead of a linear probability model, we estimate a Poisson and a Logit model and find similar (and more precisely estimated) results.

4 Mechanisms

An increase in hate crimes following larger protests could arise from three distinct mechanisms: salience, coordination, or signaling. First, larger protests may increase the salience of anti-refugee sentiment, reinforcing existing resentment and triggering a temporary rise in hate crimes as individuals are reminded of and mobilized by the issue. Second, protests can serve as coordination platforms, facilitating the organization and execution of collective acts of violence. Third, protests may function as a public signal of widespread support for anti-refugee sentiment, reducing the perceived social or legal costs of engaging in hate crimes. In extreme cases, such signaling may not only lower deterrence but also shift individual preferences by increasing the perceived social rewards or fostering a "taste" for such actions. We assess these mechanisms through a series of empirical tests designed to isolate their respective contributions.

4.1 Dynamic Effects

We begin by analyzing the dynamic treatment effects of protest size on hate crimes. Each mechanism implies a distinct temporal pattern. If salience is the primary driver, we expect an immediate spike in hate crimes followed by a rapid decline as attention to the protest fades. In contrast, a coordination mechanism would manifest as a delayed increase in hate crimes, reflecting the time needed for planning and organizing, followed by a gradual attenuation as

coordinated efforts dissipate. Finally, if signaling is the dominant mechanism, we anticipate a more persistent increase in hate crimes, stabilizing at a higher level due to a belief update about the social costs and returns to hate crimes.

Event-Study design

To investigate the short- and medium-run dynamic effects, we employ an event study design at both the municipality-day and municipality-week levels. Analyzing the dynamic treatment effects also addresses remaining concerns about identification. First, it allows us to confirm that municipalities experiencing PEGIDA protests on pleasant days were not on different pre-existing trajectories in terms of their propensity to commit hate crimes. Second, it enables us to identify inter-temporal substitution effects. If the timing of hate crimes shifts, rather than the overall propensity to commit them, this would be revealed in treatment coefficients of opposite signs during the pre- or post-protest periods. We estimate an event-study of the following form:


$$\begin{aligned}
 HC_{it} = & \sum_{\substack{k=T_0 \\ k \neq -1}}^{T_1} \beta_k \mathbb{1}(\text{weather}_{ik} \times \text{protest}_{ik}) + \gamma_k \mathbb{1}(\text{weather}_{ik}) + \phi_k \mathbb{1}(\text{protest}_{ik}) \\
 & + \alpha_i + \lambda_t + \nu_i \times T + \delta V_{it} + \epsilon_{it}
 \end{aligned} \tag{3}$$

where HC_{it} is defined as a dummy variable that switches on if any hate crime was recorded in municipality i on day or week t . Equivalent to our previous estimation, we include municipality α_i and period λ_t fixed effects, as well as municipality-specific linear time trends $\nu_i \times T$. The coefficients β_k trace the dynamic effects of PEGIDA Monday protest on pleasant days for each period k before and after pleasant day protests. The coefficients γ_k and ϕ_k trace the dynamic treatment effects of any protest or any pleasant weather day. The vector of time-varying municipal level variables V_{it} includes the full set baseline controls (but excluding the lagged controls). Again, we focus on the sample of ever-treated municipalities and cluster standard errors at the municipal level.


Persistence of hate crimes

Panel A of Figure 6 presents the results of the event study at the municipality-day level. The pre-treatment coefficients are tightly clustered around zero and statistically insignificant, indicating no differential trends in hate crimes in the days leading up to protests on pleasant days. In the post-treatment period, we observe a marked increase in hate crimes, with coefficients rising sharply immediately on the day of the protest. The effects persist over the subsequent days, with no evidence of reversals or compensatory declines in crime rates. This sustained increase rules out the hypothesis of inter-temporal substitution, and instead suggest a structural shift in hate crime behavior rather than a mere reallocation of crimes across time. Panel B displays the results aggregated to the municipality-week level. Pre-treatment coefficients remain near zero and statistically insignificant, mirroring the findings at the daily level and confirming the absence of differential pre-trends over a longer period. Post-treatment, the coefficients exhibit a

Figure 6. Event Study: Hate Crimes Following Protests on Pleasant Days



(a) Hate Crimes at the Municipality and Day Level



(b) Hate Crimes at the Municipality and Week Level

Note: Event study using `xtevent` from Freyaldenhoven et al. (2024). Regression at the municipality and day level (panel a) and municipality and week level (panel b) with linear time trends where the outcome is a dummy for any hate crime in the week of the protest, estimating equation 3. Standard errors are clustered at the municipality level. The set of controls includes all controls of Table 1 column 4, except for lagged controls, as well as a dummy for the post-protest and post-pleasant weather periods. The sample consists of ever-treated municipalities only. Vertical bars represent 90% and 95% confidence intervals, respectively.

sustained and statistically significant upward trajectory, peaking approximately 4-5 weeks after the protest and remaining at a higher equilibrium level thereafter.

Overall, the event-study results reveal that individuals who commit hate crimes do not strategically choose the timing of their actions; they do not substitute hate crimes to align with pleasant protest days. We also find no evidence that hate crimes decrease again over time. Instead, they converge to a new, higher equilibrium level, consistent with a permanent reduction in the actual or perceived costs of committing these crimes. This is consistent with other results in the literature on national election outcomes or signals from populist entrepreneurs and extend them to show that highly localized and bottom-up signals may be equally important for shifts in social norms and explain persistent geographic variation of those norms. It is also consistent with *informational cascades* (Bikhchandani et al., 1992), where shifts in norms can arise because individuals act in sequence and later individuals update their information based on the behavior of early movers. In addition, the immediate surge in hate crimes on the protest day suggests that salience, or potentially violent dynamics arising from large protests (for instance by attracting radicals from outside the municipality), may also play a role. However, the longer-run increase indicates that signaling effects, whereby large protests normalize or legitimize harmful behavior, are likely the dominant mechanism.

4.2 Spatial Diffusion

This section investigates the diffusion of hate crimes across municipalities via geographic and social media networks. Geographic proximity captures the potential for spillovers arising from physical closeness to large protests, which may amplify salience through regional media coverage or interpersonal interactions. Social media proximity, measured both within PEGIDA-specific networks and broader Twitter networks, allows us to test whether digital connectivity facilitates the spread of protest influence, either by amplifying visibility, enabling organizational efforts, or signaling normative shifts. By examining the effects of proximity-weighted exposure to protest size on hate crimes, both contemporaneously and over time, this exercise enables us to assess whether digital and physical diffusion mechanisms operate independently or complement each other. Furthermore, we test the persistence of these effects to determine whether the observed dynamics are temporary or reflect longer-term shifts in behavior.

Measures of diffusion

In a first step, we develop a social media proximity weighted measure of exposure to large protests. We define S_{it}^{Γ} as the social media proximity-weighted sum of protest participants in all other municipalities in week t , where Γ can either be the PEGIDA network P or the wider Twitter network W .

$$S_{it}^{\Gamma} = \sum_{j \neq i} N_{j \rightarrow i(t)}^{\Gamma} \times \log(1 + \text{participants})_{jt}$$

The PEGIDA social media proximity weights N^P are based on retweets of tweets that contain the word PEGIDA. Specifically, for each retweet, we locate the original user and the retweeting

user. To compute the time-varying influence from municipality j to municipality i , we count the number of retweets in municipality i coming from users in municipality j over the preceding 6 months and normalize those by the total number of active (i.e., retweeting) PEGIDA users in municipality i to account for the overall PEGIDA Twitter penetration at the extensive and intensive margin.¹⁷ It is worth noting that social media proximity can be asymmetric. In other words, the influence of municipality j on i can be larger than the influence of municipality i on j , if Twitter users in i are more likely to retweet Twitter users from j .

Complementing the PEGIDA social media proximity measure, we create the equivalent index for Twitter networks more broadly, drawing from a random sample of 600,000 tweets. Again, we take the number of (any) retweets in municipality i of users located in municipality j during the previous 6 months, and normalize by the overall retweeting activity in i . Lastly, we develop a measure for physical exposure to large protests, G_{it} , measured as the distance weighted sum of protest participants in other municipalities. Specifically, we compute the mutual influence between municipality i and municipality j as the distance between the centers of the two municipalities, applying a linearly decreasing window function until 100 km distance. Leveraging our three exposure measures and replicating our baseline specification, we estimate the following linear probability model:

$$HC_{it} = \beta_1 S_{it}^P + \beta_2 S_{it}^W + \beta_3 G_{it} + \beta_4 P_{it} + \gamma X_{it} + \delta_t + \mu_{im} + \nu_i \times t + \epsilon_{it} \quad (4)$$

We are particularly interested in the coefficient β_1 for both contemporaneous and subsequent hate crimes. A large and persistent effect of PEGIDA social media networks on hate crimes, in our view, supports the signaling hypothesis since these networks amplify and sustain the normative shifts initiated by large protests. By broadcasting messages and imagery that portray anti-refugee sentiment as widely supported, PEGIDA networks could embolden individuals already sympathetic to the cause.

We face three challenges in isolating the effect of β_1 on hate crimes. First, PEGIDA social media networks could just be a proxy for geographic proximity or overall digital connectedness. If there are spatial spillovers of protests and hate crimes, we may falsely attribute this to social media networks. Including both geographic proximity and wider social media networks allows us to distinguish right-wing networks from other forms of exposure. It also allows us to distinguish between audiences, that is large protest as a signal to the broader public versus large protest as a signal to a sympathetic users, i.e. PEGIDA re-tweeters.

Second, we face a simultaneity problem. Digitally connected municipalities could experience a surge in protest attendance that is inspired by protest attendance elsewhere. In this case, we would not strictly identify the effect of social media networks but that of local protest attendance driven by social media networks. This is unlikely because we measure protest on the same day and protest mobilization is unlikely to happen in real time. Nevertheless, we address this concern in two ways: for one, we can directly measure the effect of social media networks on

¹⁷Our scraping and geo-location method is described in more detail in Appendix C. Note that we can only geo-locate tweets at the NUTS-3 level, a higher level of geographic aggregation. We assign the same weights to all municipalities within the same NUTS-3 region.

protest attendance. In addition, when looking at the effect on hate crimes, we always condition on local protest size and therefore measure the additional effect of being connected to places with large protests.

Lastly, digitally connected municipalities may react to the same shock that drives hate crimes and protest participation. We address this concern by investigating the persistence of the effect. We look at the effect of S_{it}^{Γ} on hate crimes in t as well as $t + 1$. If radicals update their beliefs about the preferences of others, this effect should persist over time. If digitally connected municipalities react to the same time-varying factor, we may not observe the same level of persistence. Importantly, we simultaneously control for local protest attendance in t , as well as past protest, past protest attendance and hate crimes in $t - 1$. This means that we capture deviations from previous levels, isolating the additional marginal effect of protest attendance in digitally connected municipalities.

Spread of hate crimes through geographic and social media networks

Table 2 presents our results. Columns 1 to 3 successively include PEGIDA social media proximity, wider social media proximity, as well as geographic proximity to large protest locations before including all variables in a horse race in column 4. Column 5 investigates the persistence of the effect by looking at the likelihood of observing a hate-crime in the following week. The last two columns, column 6 and 7, address the issue of simultaneity by looking at the direct effect of our three exposure measures on contemporaneous protest occurrence and size.

In a first step, we focus on our coefficient of interest β_1 . In columns 1 and 4, the coefficients on S_{it}^P are both large and highly significant, indicating that municipalities connected to areas with large protests via PEGIDA networks experience higher rates of hate crimes, even after controlling for local protest attendance (column 1) and wider social media networks and geographic proximity (column 4). This effect persists into the following week (column 5), suggesting that these networks amplify and sustain the normative shifts initiated by large protests.

In contrast, the effects of broader social media proximity are smaller and less persistent. While S_{it}^W significantly predicts hate crimes in the same week (column 2), the effect diminishes by the following week (column 5). This pattern aligns with the salience mechanism, where broader digital networks amplify the visibility of protests temporarily but do not sustain the normative shifts needed for longer-term behavioral changes. The results for geographic proximity G_{it} are significant but small, indicating a more localized diffusion effect. The lack of persistence and the smaller magnitude relative to S_{it}^P highlight the limited role of physical proximity, which (similar to the wider social media network) serves as a signal to a broader population rather than a group that already engages with and is interested in PEGIDA.¹⁸

Importantly, columns 6 and 7 show no significant effects of S_{it}^P , S_{it}^W , or G_{it} on local protest occurrence or protest size, suggesting that these proximity measures influence hate crimes independently of additional protest mobilization. Importantly, this finding not only mitigates concerns about simultaneity but it also addresses concerns related to substitution or spill-over effects across space. Reassuringly, we find no evidence that proximity to large protest locations

¹⁸We show in Table ?? that PEGIDA related tweets are highly correlated with a more direct measure of support: the number of followers of the official PEGIDA account.

leads to a lower likelihood of recording protests or that it leads to a decrease in protest size elsewhere, ruling out the possibility that hate crimes simply shift from control municipalities to treatment municipalities.¹⁹

Complementing work that shows how social media enables the rapid diffusion of *ad hoc* protests (Qin et al., 2024), our context involves pre-announced demonstrations that do not spread across the network. Instead, the aftermath of larger protests appears to reverberate through online connections, raising hate crime even absent protest events themselves. Finally, our findings underscore that not all online networks function alike: whereas broader social media engagement has a smaller and shorter-lived effect, right-wing networks composed of ideologically sympathetic users propagate the norms that condone radical action.

¹⁹In addition, this exercise mirrors and extends the robustness checks that concern reporting bias that may be driven by an increase in police effort. If police resources are directed to locations with large protests, but not to locations that are more connected through PEGIDA Twitter networks (conditional on geographic distance and overall Twitter connectedness), the coefficient will not capture bias in crime reporting.

Table 2. Diffusion: geographic and (right-wing) social media proximity to other large protest locations

	Hate crime in t				Hate crime in t + 1	Protest in t	Log(participants) in t
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log(participants)	0.0216*** (0.00515)	0.0207*** (0.00496)	0.0215*** (0.00501)	0.0208*** (0.00494)	-0.000240 (0.00526)		
PEGIDA social media prox. S_{it}^P	0.00238*** (0.000307)			0.00220*** (0.000305)	0.00225*** (0.000306)	0.000408 (0.000280)	0.00209 (0.00152)
Overall social media prox. S_{it}^W		0.00142*** (0.000277)		0.000444** (0.000173)	0.000300 (0.000200)	-8.32e-05 (0.000229)	-0.000502 (0.00103)
Geographic prox. G_{it}			0.000174*** (6.46e-05)	0.000116* (6.47e-05)	5.15e-05 (6.26e-05)	-3.92e-05 (3.30e-05)	-0.000204 (0.000163)
Observations	2,771,167	2,771,167	2,771,167	2,771,167	2,771,167	2,771,167	2,771,167
Municipalities	10,677	10,677	10,677	10,677	10,677	10,677	10,677
Adj. R-squared	0.092	0.093	0.093	0.093	0.093	0.337	0.310
Mean dep. var.	0.00282	0.00282	0.00282	0.00283	0.00283	0.00282	0.00282
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality × MoY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality linear time trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes
All controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Differences in differences estimation with week and municipality-month of year fixed effects and municipality linear time trends. Full sample of municipalities. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Controls include GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants respectively, all measured at the end of each calendar year. S_{it} is the social media network proximity-weighted sum of protest participants in other municipalities (standardized). Superscript P indicates PEGIDA network, superscript W indicates wider network, based on random sample of tweets. G_{it} geographic proximity weighted (linearly decreasing window function until 100km) sum of participants in other municipalities (standardized). Social media proximity is measured as the per capita number of retweets in municipality i over the previous six months of sampled tweets coming from municipality j . Outcome in columns 1 to 4 is any hate crime in the same week. Outcome in column 5 is any hate crime in the following week. Outcomes in columns 6 and 7 are whether a protest occurred on the same day and the log number of participants, respectively.

4.3 Counter-Mobilization

We next ask whether large protests spark a response from opponents of the movement, who may similarly update their beliefs about the relative benefits of mobilization. In principle, pro-refugee activists could respond either on the streets (i.e., by organizing counter-protests) or online (e.g., by voicing support for refugee issues on social media).

To examine physical counter-protests, we collect newspaper articles referencing PEGIDA from the GENIOS database, which compiles coverage from 282 publications between 2015 and 2020. Using a customized large language model, we identify articles that mention counter-demonstrations and extract their location, date, and reported size. Appendix C reports more details on this. Columns 1 to 6 of Table 3 repeat our baseline 2SLS framework with six different outcomes: the occurrence, frequency, and size of counter-protests in both the same and the subsequent week. Throughout, we find no evidence that larger PEGIDA protest generate protests in opposition to it. One possibility is that local activists prefer not to confront potentially hostile crowds, or that municipalities hosting PEGIDA lack a critical mass of motivated pro-refugee supporters.

To investigate whether activists rely more heavily on safer or less effortful online mobilization, we scrape all German-language tweets containing the hashtag `#refugeesWelcome` from 2015 to 2018.²⁰ This hashtag represents a unifying banner for pro-refugee advocates. Using machine-learning sentiment analysis (Guhr et al., 2020) and a manual classification of 1,000 tweets, we verify that tweets carrying the hashtag indeed convey positive attitudes toward refugees. We also gather the universe of PEGIDA-related tweets for our entire sample period and apply the same sentiment classification. In columns 7 and 8, we regress, respectively, the number of pro-refugee tweets in a municipality-week and the net positive sentiment among PEGIDA tweets (i.e., positive minus negative over the total count) on protest size. Although the estimates suggest a potentially moderate increase in pro-refugee activity online and somewhat diminished positive sentiment toward PEGIDA, both effects are imprecise. Thus, we find at most a muted counter-mobilization on digital platforms. Taken together, these results suggest that the main impact of bigger PEGIDA protests is not met by sizable or immediate counter-protests offline, and any shift to online activism is large in magnitude but noisily estimated.

²⁰We cannot extend this dataset to 2020 due to Twitter API limitations.

Table 3. **Large Protests and Counter-Mobilization Offline and Online**

	Counter-Protest						Pro-Refugee Twitter	
	occurrence		frequency		size		#refugeeswelcome	PEGIDA sentiment
	t	t+1	t	t+1	t	t+1	t	t
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(participants)	0.0417 (0.246)	-0.0559 (0.0877)	0.414 (0.402)	-0.0629 (0.264)	1.677 (1.122)	-0.799 (0.759)	6.538* (3.913)	-1.215 (2.191)
Observations	21,678	21,595	21,678	21,595	21,678	21,595	17,362	19,914
Municipalities	84	84	84	84	84	84	84	84
Mean dep. var.	0.0280	0.0273	0.0496	0.0469	0.133	0.128	1.520	-0.0956
F first	10.79	10.92	10.79	10.92	10.79	10.92	7.012	12.99
All controls & FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: TWFE 2SLS estimation with week and municipality-month of year fixed effects and municipality linear time trends. Ever-treated sample only. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Instrument is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleinbergen-Paap F-Statistics are reported at the bottom of the table. Controls comprise GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants, all measured at the end of each calendar year; as well as the interaction between controls and the weather as well as the protest dummy. Outcomes for “occurrence,” “frequency,” and “size” in columns (1) to (6) capture whether a counter-protest occurred, the number of such events, and their total reported attendance, in the same week (including Monday) and the following week respectively, using data on counter-protests from the newspaper archive GESIOS. Column (7) measures the total #refugeeswelcome tweets in the municipality-week, while column (8) reports net positive sentiment in PEGIDA-related tweets (positive minus negative, divided by all PEGIDA tweets).

4.4 Media coverage

We now examine whether local exposure to PEGIDA-related newspaper articles, in conjunction with broader social media exposure to PEGIDA-related tweets, influences hate crimes over and above the direct effect of local protest participation. This analysis extends our analyses on spatial diffusion and counter-mobilization in the previous sections by showing how traditional newspaper and social media coverage can reinforce or counter right-wing messaging. Importantly, we now broaden our lens to capture both *pro*- and *anti*-PEGIDA protest coverage in newspapers, as well as PEGIDA-related and #refugeesWelcome content on social media. This enables us to gauge whether certain types of coverage and online activity amplify protest-driven hate crimes or instead mitigate them by offering a countervailing message.

Measures of media exposure

We construct a weekly newspaper-exposure measure based on the GENIOS database that contains the universe of newspaper articles in Germany, matched to municipality-level readership from advertising statistics (see Appendix C for more detail). For each article referencing pro-PEGIDA or anti-PEGIDA demonstrations, we multiply the newspaper’s circulation in each municipality by the number of published articles in the previous week, then sum over all newspaper outlets and divide by the municipality’s population to capture per capita exposure. Similarly, we incorporate municipality-level social media exposure both to tweets supportive of PEGIDA and to tweets using the pro-refugee hashtag #refugeesWelcome. We take the sum of tweets in the previous week and multiply this with the number of Twitter users in the municipality from a random sample of tweets, divide this by the overall population and standardize for comparability. Our exposure measures write as follows, where $\sigma \in (\text{pro}, \text{anti})$ suggests pro and anti-PEGIDA coverage and p is an indicator for one of 286 newspaper outlets in Germany:

$$\text{Newspaper Exposure}_{m,t}^{\sigma} = \frac{\sum_p \text{Articles PEGIDA}_{p,t-1}^{\sigma} \times \text{Readership}_{p,m}}{\text{Population}_m}$$
$$\text{Twitter Exposure}_{m,t}^{\sigma} = \frac{\text{Tweets PEGIDA}_{t-1}^{\sigma} \times \text{Twitter Users}_m}{\text{Population}_m}$$

We estimate our standard linear probability model, conditioning on local protest size and examine the effect of newspaper and Twitter exposure to pro or anti-PEGIDA content. Note that our estimation follows the logic of a shift-share instrument, the shifter being the number of articles or tweets in the previous week, and the share being the respective audience at the local level. We control for local protest participation to account for the fact that local protest size will also determine its coverage. The coefficients for the exposure measures will thus identify the effect of local media exposure above and beyond local attendance.

Amplification of hate crimes through newspaper and social media coverage

Table 4 reports our results. Columns 1 to 3 regress the probability of at least one hate crime in the current week on protest size and various coverage variables. Column 4 studies the persistence

of the effect. Throughout, we focus on the sample of ever-treated municipalities and condition on the size of the local protest and the same set of fixed effects and controls as in our baseline specification.

In a first step, we examine the relationship between exposure to pro- and anti-PEGIDA content on Twitter. The positive and statistically significant coefficient on pro-PEGIDA tweets indicates that online messages supportive of the movement bolster hate crimes above and beyond local protest attendance. In contrast, pro-refugee tweets do not counter-balance this effect. Next, we repeat this exercise and examine the newspaper coverage of protest in support of and against PEGIDA. Municipalities with a larger readership of newspapers that cover PEGIDA protest in the previous week, experience a surge in hate crimes above and beyond local protest participation. At the same time, anti-PEGIDA newspaper exposure exerts a weak negative effect (albeit significantly smaller in magnitude), suggesting that print coverage of counter-demonstrations may slightly mitigate radical action.

Next, column 3 includes all measures simultaneously. All coefficients remain remarkably similar, suggesting that social media and newspaper exposure are not strongly correlated. Pro-PEGIDA content on social media and in newspapers maintains a significant and positive effect on the likelihood of observing a hate crime in the same week, while the negative effect of newspaper coverage of anti-PEGIDA events becomes noisy. The coefficient for Twitter exposure is substantially larger than that of newspaper coverage. This could be for two reasons. First, social media users are more likely to be at the margin of engaging in radical action as compared to the newspaper readership. Second, the type and tonality of coverage of PEGIDA in newspapers could have a moderating effect while social media coverage might be more inciting. Finally, column 4 examines the persistence of the effect. Mirroring our previous findings in section 4.2, we find that social media exposure has the strongest and most persistent effect on radical action.

Overall, these findings reinforce that media ecosystem boost the signal offered by PEGIDA protests, above and beyond the possibility to coordinate locally. Pro-PEGIDA coverage on both traditional and social media boosts hate crimes contemporaneously and, in the case of social media, into the following week. Anti-PEGIDA content exerts at most a minor restraining effect, indicating that strong, favorable messaging can overshadow weaker or less persistent counter-mobilization. Although past research suggests that counter-protests can lower right-wing vote shares in elections (Lagios et al., 2025), our findings here, combined with the negligible counter-mobilization in response to large PEGIDA protests, imply that there is no analogous symmetry of opposition and no consequent reduction in hate crimes.

4.5 Coordination

This section investigates whether hate crimes after large right-wing protests arise through coordinated efforts or through individual-level actions spurred by a broader sense of impunity. We begin by analyzing whether perpetrators are more likely to have extremist backgrounds, act in groups, or commit offenses in a public manner. We then consider football matches as a placebo to assess whether their potential for gathering and mobilization (especially among far-right hooligan fan bases) translates into a similar uptick in hate crimes.

Table 4. **Newspaper and Social Media Coverage pro- & anti-PEGIDA**

	Hate crime in t			Hate crime in t + 1
	(1)	(2)	(3)	(4)
Log(participants)	0.0173*** (0.00516)	0.0193*** (0.00493)	0.0174*** (0.00516)	-0.00216 (0.00432)
Twitter Exposure ^{pro}	1.514** (0.591)		1.521** (0.594)	1.174** (0.476)
Twitter Exposure ^{anti}	0.372 (0.277)		0.375 (0.279)	0.287 (0.434)
Newspaper Exposure ^{pro}		0.0477*** (0.0166)	0.0427** (0.0169)	-0.00172 (0.0125)
Newspaper Exposure ^{anti}		-0.0332* (0.0199)	-0.0213 (0.0190)	0.0203 (0.0212)
Observations	17,445	21,678	17,445	17,445
Municipalities	84	84	84	84
Mean dep. var.	0.0938	0.0915	0.0938	0.0938
All FE & controls	Yes	Yes	Yes	Yes

Note: Differences in differences estimation with week and municipality-month of year fixed effects and municipality linear time trends. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Controls include GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants respectively, all measured at the end of each calendar year. Treatment is defined as in section 4.4. Outcome in columns 1 to 3 is any hate crime in the same week. Outcome in column 4 is any hate crime in the following week.

Perpetrator and crime characteristics

We use a large language model (LLM) to classify detailed textual descriptions of hate crimes into multiple, thematically grouped indicators such as perpetrator background, type of offense, and location. Concretely, we fine-tune a ChatGPT model to parse raw text and return a Python data frame that places each piece of information into predefined categories, including whether the perpetrator was already known to law enforcement, acted in a group or alone, spontaneously, in a public space or whether hate crimes were committed in relation to a protest or rally. We describe this procedure in more detail in Appendix C.

Table 5 builds on these classifications to examine the effect of larger right-wing protests on key perpetrator characteristics and crime types. Our outcome in column 1 captures whether crimes are committed by individuals with a documented extremist or recidivist record; the estimates show a clear increase in these cases, suggesting that perpetrators already prone to radical acts feel emboldened after larger protests. Next, column 2 focuses on group-based offenses. The estimate is small in magnitude and insignificant, indicating that large protests do not translate into more collective attacks. In contrast, column 3 examines hate crimes with perpetrators that act alone, and the findings reveal a notable rise in single-offender incidents, consistent with the idea that lowered social stigma enables individuals to commit hate crimes without the safety of group support or need to coordinate. Column 4 tracks the presence of unplanned or impulsive acts, but there is no apparent change in spontaneous crimes, which is consistent with our previous findings in section 4.1 and suggests that these incidents reflect enduring shifts in rather than short-lived emotional outbursts. Column 5 focuses on whether

offenses occur in open, heavily trafficked venues (public squares, public transport, train stations etc.); the significant increase in that share points to a decline in perpetrators’ fear of social disapproval, as they are now more willing to engage in visible displays of aggression. Finally, column 6 captures protest-related offenses. We find no significant increase, indicating that the surge in hate crimes does not revolve around specific protest events. Combined these results suggest that the surge in hate crimes stems from radicalized individuals feeling individually licensed to offend.

Table 5. Effect of Larger Protests on Hate-Crime Characteristics

	known perp. (1)	group perp. (2)	single perp. (3)	spontaneous (4)	public space (5)	protest related (6)
Log(participants)	0.513** (0.215)	0.0195 (0.164)	0.459* (0.237)	0.262 (0.185)	0.326** (0.163)	0.186 (0.131)
Observations	21678	21678	21678	21678	21678	21678
F first	10.77	10.77	10.77	10.77	10.77	10.77
Municipalities	84	84	84	84	84	84
Mean dep. var.	0.0915	0.0443	0.0375	0.0688	0.0186	0.00738
All controls & FEs	Yes	Yes	Yes	Yes	Yes	Yes

Note: TWFE 2SLS estimation with week and municipality-month of year fixed effects and municipality linear time trends. Ever-treated sample only. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Instrument is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleibergen-Paap F-Statistics are reported at the bottom of the table. Controls comprise GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants, all measured at the end of each calendar year; as well as the interaction between controls and the weather as well as the protest dummy. Outcomes categorized by LLM based on descriptions of over 9,000 hate crimes. Column (1) includes hate crimes where the perpetrator was known to the police, has committed hate crimes before, or belonged to a know right-wing extremist group. Columns (2) and (3) identify whether the hate crimes was committed alone or in a group. Columns (4) to (6) examine whether the hate crime happened spontaneously, whether it happened in a frequented space (like public transport, public squares or train stations) and whether the hate crime happened in the context of a protest or rally.

Opportunities for coordination

To further distinguish the impact of protest participation as a public signal from its role as a coordination device for far-right supporters, we use football matches as a placebo check. If coordination constraints are binding, football matches should increase hate-crimes, particularly when the fan-base is sympathetic to right-wing ideologies.

To test this, we collected data on football match outcomes from the website fbref.com, which offers comprehensive statistics for various domestic and international football leagues. For Germany, data is available for three domestic male leagues as well as the national cup (*Bundesliga*, *2. Bundesliga*, *3. Liga*, and *DFB-Pokal*, respectively). Through web-scraping, we assemble a data set for all seasons spanning from 2014 to 2020, including match date and time, participating teams, scores, attendance, venue name, and additional unstructured notes such as whether the game ended with extra time or penalties. We complement this data with information on the political ideology of football clubs’ fan-base. Our primary source is Duben (2015) who highlights 9 clubs as having far-right fanbases. We complement this information with a report from the Federal Agency for Civic Education (Claus, 2024) and with a recent New York Times article by Hughes (2024). We verify the classification against the teams mentioned in relevant forum discussions on the social media platform Reddit.com. We present more details

on the data collection and classification method in Appendix C.

Table A.4 examines the impact of football matches on hate crimes against refugees. The analysis mirrors the previous two-way fixed effects model but now leverages municipality and day variation.²¹ The municipality by day of the week fixed effect absorbs any time-invariant municipality specifics that may vary for each day of the week. If a municipality has local factors that affect their level of hate crimes and certain weekly patterns may further exacerbate them (Cohn & Rotton, 2003), this fixed effect will absorb the systematic differences. Day fixed effects absorb any common shocks or seasonal trends affecting all municipalities equally on a given date. All specifications also control for the number of hate crimes and football matches on the previous day as well as accounting for long-term linear trends within municipalities.

In columns 1 to 3, we examine the effect of the occurrence and size of any football match on the likelihood of observing a hate-crime on the same day, successively adding fixed effects and controls. Our estimated coefficients are negative and not significantly different from zero.²² Next, in columns 4 to 7, we explore the possibility that hate crimes are exacerbated on days with contested matches or by supporters who are classified as right-wing hooligans. If hate crimes are driven by coordination then the effects should be particularly strong on days when a football team with hooligan fans plays. If hate crimes arise spontaneously due to heightened emotions and frustration then a derby or a contested match (as measured by penalties or extra time) should drive an increase. Throughout, there is no observable effect on hate crimes following either matches with hooligan fans or emotionally loaded matches.

5 Conclusion

Our study shows that broader participation in far-right protests, instrumented by weather-induced variation, substantially increases hate crimes in both the short and medium run. Leveraging high-resolution data and a rich set of robustness checks, we identify a 0.5 percentage point rise in the probability of a hate crime for each 1% increase in protest size. This link appears driven less by direct coordination among participants than by a signal about wider support for anti-refugee sentiment. Put simply, when a protest attracts more marginal supporters, it shifts local norms around radical action, emboldening individuals—especially those predisposed to extremism—to engage in hate crimes.

Our findings carry three main implications. First, they underscore that public events, even when ostensibly peaceful, can embolden extremist behavior. That mechanism suggests that policy proposals aimed at banning or restricting such protests may be viewed as attempts to safeguard vulnerable groups, but also confront serious trade-offs. Limiting public assemblies raises concerns about free speech and risks heightening a sense of victimization among protesters. Policymakers thus face a nuanced challenge: interventions must weigh the right to protest against the tangible harm that protest signals can produce. Second, the amplification via social

²¹We do this because we do not have one-off events that always take place on the same day of the week.

²²The literature on football hooliganism in Germany (Andres et al., 2023) suggests that the majority of violent assaults in response to a football match occur on the day of the match. However, it is possible that coordination takes time. Hence, in Table A.5, we examine the effect of football matches over the following one to seven days, as well as the number of hate crimes or the type of hate crime. Similar to our previous findings, there is no systematic relationship between matches and hate crimes.

media and local media coverage highlights that governments, law enforcement, and civil society cannot focus solely on events “on the ground.” Digital channels accelerate and sustain the norm-shifting impact of mass protest, suggesting that initiatives to monitor extremist content online or counter misinformation may help mitigate downstream hate crimes.

Finally, while our evidence pinpoints the mechanisms of signaling, salience, and low-cost diffusion through social networks, several open questions remain. Future work could investigate potential feedback loops between protest attendance, policy changes and political extremism. PEGIDA arguably led to the success of the right-wing party AfD in Germany, functioning as a bridge to civil society and thus bringing extremist views into the mainstream. It may also prove fruitful to study whether distinct forms of counter-speech or bystander interventions offset the negative externalities we document - particularly in contexts where counter-protests struggle to take hold. We believe these avenues would enrich our understanding of how collective action, media ecosystems, and social norms interact to shape extremist behavior.

Taken together, our results indicate that even minor expansions in protest attendance carry significant consequences for public safety, especially when they embolden those already prone to violence. By uncovering the mechanisms through which seemingly moderate participants inadvertently shift the norms of radical action, this paper highlights how protest dynamics can alter the political landscape far beyond the ballot box.

References

- Acemoglu, D. & Robinson, J. A. (2000). Why did the west extend the franchise? democracy, inequality, and growth in historical perspective. *The Quarterly Journal of Economics*, 115(4), 1167–1199.
- ACLEDD (2024). Aclcd conflict index results: July 2024. <https://acleddata.com/conflict-index/index-july-2024/>. Accessed: 2025-02-06.
- Ajzenman, N., Cavalcanti, T., & Da Mata, D. (2023). More than words: Leaders' speech and risky behavior during a pandemic. *American Economic Journal: Economic Policy*, 15(3), 351–371.
- Andreoni, J. & Bernheim, B. D. (2009). Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, 77(5), 1607–1636.
- Andres, L., Fabel, M., & Rainer, H. (2023). How much violence does football hooliganism cause? *Journal of Public Economics*, 225, 104970.
- Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2), 169–217.
- Beraja, M., Kao, A., Yang, D. Y., & Yuchtman, N. (2023). AI-tocracy*. *The Quarterly Journal of Economics*, 138(3), 1349–1402.
- Bernhardt, A., Field, E., Pande, R., Rigol, N., Schaner, S., & Troyer-Moore, C. (2018). Male social status and women's work. In *AEA Papers and Proceedings*, volume 108 (pp. 363–367).
- Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, 102(5), 841–877.
- Berntzen, L. E. & Weisskircher, M. (2016). Anti-islamic pegida beyond germany: Explaining differences in mobilisation. *Journal of Intercultural Studies*, 37(6), 556–573.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992–1026.
- Boyer, P., Gauthier, G., Yaouanq, Y. L., Rollet, V., & Schmutz, B. (2024). The lifecycle of protests in the digital age. *CESifo Working Paper*, (No. 11257).
- Bray, K., Braakmann, N., & Wildman, J. (2022). Austerity, welfare cuts and hate crime: Evidence from the UK's age of austerity. *Journal of Urban Economics*, (pp. 103439).
- Bursztyjn, L., Egorov, G., & Fiorin, S. (2020). From extreme to mainstream: The erosion of social norms. *American Economic Review*, 110(11), 3522–3548.
- Bursztyjn, L. & Yang, D. (2022). Misperceptions about others. *Annual Review of Economics*, 14, 425–452.
- Cantoni, D., Kao, A., Yang, D. Y., & Yuchtman, N. (2024). Protests. *Annual Review of Economics*, 16.
- Carr, J., Clifton-Sprigg, J., James, J., & Vujic, S. (2020). *Love Thy Neighbour? Brexit and Hate Crime*. IZA Discussion Paper DP No. 13902.
- Casanueva, A. (2021). Can chants in the street change parliament's tune? the effects of the 15M social movement on spanish elections. *mimeo*.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1), C1–C68.

- Chersich, M. F., Swift, C. P., Edelstein, I., Breetzke, G., Scorgie, F., Schutte, F., & Wright, C. Y. (2019). Violence in hot weather: Will climate change exacerbate rates of violence in South Africa? *South African Medical Journal*, 109(7), 447–449.
- Claus, R. (2024). *Extrem rechte Fußballfans und die Nationalmannschaft des DFB*. Technical report, Bundeszentrale für politische Bildung.
- Cohn, E. & Rotton, J. (2003). Even criminals take a holiday: Instrumental and expressive crimes on major and minor holidays. *Journal of Criminal Justice*, 31, 351–360.
- Conley, T. G. (1999). Gmm estimation with cross sectional dependence. *Journal of Econometrics*, 92(1), 1–45.
- Correa, S., Nandong, G., & Shadmehr, M. (2025). Grievance shocks and coordination in protest. *American Journal of Political Science*, 69(1), 341–354.
- Correia, S., Guimarães, P., & Zylkin, T. (2019). Ppmlhdf: Fast poisson estimation with high-dimensional fixed effects. *arXiv preprint arXiv:1903.01690*, 13.
- De Chaisemartin, C. & d’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–96.
- De Chaisemartin, C. & D’Haultfoeuille, X. (2022). *Difference-in-Differences Estimators of Intertemporal Treatment Effects*. Working Paper 29873, National Bureau of Economic Research.
- De Chaisemartin, C. & D’Haultfoeuille, X. (2022). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey. *The Econometrics Journal*, 26(3), C1–C30.
- Dipoppa, G., Grossman, G., & Zonszein, S. (2023). Locked down, lashing out: Covid-19 effects on asian hate crimes in italy. *The Journal of Politics*, 85(2), 389–404.
- Duben, D. (2015). *Strategien gegen Rechtsextremismus im Fußballstadion*.
- Entorf, H. & Lange, M. (2019). Refugees welcome? Understanding the regional heterogeneity of anti-foreigner hate crimes in Germany. *ZEW-Centre for European Economic Research Discussion Paper*, (19-005).
- Field, S. (1992). The effect of temperature on crime. *The British Journal of Criminology*, 32(3), 340–351.
- Freyaldenhoven, S., Hansen, C. B., Pérez Pérez, J., Shapiro, J. M., & Carreto, C. (2024). *Policy effect estimation and visualization in linear panel event-study designs: Introducing the xtevent package*. Technical report, Working Papers.
- Goetzmann, W. N., Kim, D., Kumar, A., & Wang, Q. (2015). Weather-induced mood, institutional investors, and stock returns. *The Review of Financial Studies*, 28(1), 73–111.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
- Grosjean, P., Masera, F., & Yousaf, H. (2022). Inflammatory Political Campaigns and Racial Bias in Policing*. *The Quarterly Journal of Economics*, 138(1), 413–463.


- Guhr, O., Schumann, A.-K., Bahrmann, F., & Böhme, H. J. (2020). Training a broad-coverage German sentiment classification model for dialog systems. In N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk, & S. Piperidis (Eds.), *Proceedings of the Twelfth Language Resources and Evaluation Conference* (pp. 1627–1632). Marseille, France: European Language Resources Association.
- Han, S., Riddell, J. R., & Piquero, A. R. (2023). Anti-Asian American Hate Crimes Spike During the Early Stages of the COVID-19 Pandemic. *Journal of Interpersonal Violence*, 38(3-4), 3513–3533. PMID: 35657278.
- Hanes, E. & Machin, S. (2014). Hate crime in the wake of terror attacks: Evidence from 7/7 and 9/11. *Journal of Contemporary Criminal Justice*, 30(3), 247–267.
- Heilmann, K., Kahn, M. E., & Tang, C. K. (2021). The urban crime and heat gradient in high and low poverty areas. *Journal of Public Economics*, 197, 104408.
- Hughes, S. (2024). 'If we see that flag again, we'll shoot you': German football and the far right. *The New York Times*.
- Jaschke, P., Sardoschau, S., & Tabellini, M. (2022). *Scared Straight? Threat and Assimilation of Refugees in Germany*. Technical report, National Bureau of Economic Research.
- Jiang, J., Murrugara-Llerena, N., Bos, M. W., Liu, Y., Shah, N., Neves, L., & Barbieri, F. (2022). Sunshine with a chance of smiles: How does weather impact sentiment on social media? In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16 (pp. 393–404).
- Kandel, E. & Lazear, E. P. (1992). Peer pressure and partnerships. *Journal of Political Economy*, 100(4), 801–817.
- Kanol, E. & Knoesel, J. (2021). Right-wing extremist mobilization in germany. *WZB Berlin Social Science Center, Version*, 1(0), 125–41.
- Kuran, T. (1989). Sparks and prairie fires: a theory of unanticipated political revolution. *Public Choice*, 61(1), 41–74.
- Lagios, N., Méon, P.-G., & Tojerow, I. (2025). Media, spillovers, and social norms: The electoral impact of anti-far right protests in the 2002 french election. *The Economic Journal*, (pp. ueaf001).
- Larreboure, M. & González, F. (2021). The impact of the women's march on the us house election. *Pontificia Universidad Católica de Chile, (Maret 2021)*, (pp. 4–5).
- Levy, R. & Mattsson, M. (2023). The effects of social movements: Evidence from# metoo. *Available at SSRN 3496903*.
- Lin, W. & Wooldridge, J. M. (2019). Testing and correcting for endogeneity in nonlinear unobserved effects models. In *Panel Data Econometrics* (pp. 21–43). Elsevier.
- Madestam, A., Shoag, D., Veuger, S., & Yanagizawa-Drott, D. (2013). Do political protests matter? evidence from the tea party movement. *The Quarterly Journal of Economics*, 128(4), 1633–1685.
- Mellon, J. (2021). Rain, rain, go away: 176 potential exclusion-restriction violations for studies using weather as an instrumental variable. *Available at SSRN 3715610*.
- Müller, K. & Schwarz, C. (2021). Fanning the flames of hate: Social media and hate crime. *Journal of the European Economic Association*, 19(4), 2131–2167.

- Müller, K. & Schwarz, C. (2023). From hashtag to hate crime: Twitter and antiminority sentiment. *American Economic Journal: Applied Economics*, 15(3), 270–312.
- Passarelli, F. & Tabellini, G. (2017). Emotions and political unrest. *Journal of Political Economy*, 125(3), 903–946.
- Perez-Truglia, R. & Cruces, G. (2017). Partisan interactions: Evidence from a field experiment in the United States. *Journal of Political Economy*, 125(4), 1208–1243.
- Qin, B., Strömberg, D., & Wu, Y. (2024). Social Media and Collective Action in China. *Econometrica*, 92(6), 1993–2026.
- Romarri, A. (2020). *Do Far-Right Mayors Increase the Probability of Hate Crimes? Evidence from Italy*. Working Paper 3506811, SSRN. Accessed: 2025-02-18.
- Roth, J., Sant’Anna, P. H., Bilinski, A., & Poe, J. (2022). What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature. *arXiv preprint arXiv:2201.01194*.
- Sarsons, H. (2015). Rainfall and conflict: A cautionary tale. *Journal of development Economics*, 115, 62–72.
- Sharoff, S. (2006). Open-source corpora: Using the net to fish for linguistic data. *International Journal of Corpus Linguistics*, 11(4), 435–462.
- Vorländer, H., Herold, M., & Schäller, S. (2018). *PEGIDA and new right-wing populism in Germany*. Springer.
- Wooldridge, J. (2021). Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators. *Available at SSRN 3906345*.
- Zhong, B. & Zhou, Y. (2012). “Under the weather”: The weather effects on US newspaper coverage of the 2008 Beijing Olympics. *Mass Communication and Society*, 15(4), 559–577.

Online Appendix


Appendix A: Additional Results

Figure A.1. Number and Size of PEGIDA protests over weekdays (2015-2020)




Note: left panel shows the number of protests organized by PEGIDA and their local off-shoots between 2015 and 2020 for each day of the week. The right panel displays the cumulative number of participants in protests organized by PEGIDA and their local off-shoots between 2015 and 2020 for each day of the week.

Figure A.2. Protest participation and weather cut-offs




Note: Coefficient plot for the interaction between protest day and various weather cut-offs. We estimate the following regression: $\text{Log}(1 + \text{participants})_{it} = \beta \text{Protest}_{it} \times \text{Weather Cut}_{it} + \eta_1 \text{Protest}_{it} \times X'_{it} + \eta_2 \text{Weather Cut}_{it} \times X'_{it} + \mu_{im} + \delta_t + \tau T_i + \epsilon_{it}$. We estimate separate regressions for each cut-off variable, comparing participation within this weather cut-off to participation within all other cut-offs. Dependent variable is the log of 1+ number of participants. We include municipality month of the year fixed effects μ_{im} as well as week fixed effects δ_t . We also include the weather cut-off variable and the protest variable interacted with the full set of controls ($\text{Weather Cut}_{it} \times X'_{it}$ and $\text{Protest}_{it} \times X'_{it}$). Standard errors are clustered at the municipality level. 95 percent confidence intervals are reported. Left panel uses precipitation cut-offs. We create a dummy variable that switches on if the maximum precipitation between 2pm and 5pm lies between 0-5 mm, 5-10mm, 10-15mm, > 15mm and run separate regressions with each cut-off as the Weather Cut_{it} variable. Right panel shows cut-offs for average temperature on the protest day between 2pm and 5pm for 0-7, 7-14, 14-21, 21-28 degrees Celsius. Horizontal lines represent 90% and 95% confidence intervals.

Figure A.3. Baseline result with alternative weather instruments



Note: Coefficient plot for 2SLS estimation of baseline regression using various pleasant weather cut-offs. Baseline cutoffs are 10 mm for rain and 22 Celsius for max temperature and 0 Celsius for min temperature. First three coefficients reduce rain cutoff by 1 to 3mm, 4th to 7th coefficients reduce max temperature by 1 to 3 degrees Celsius, 8th to 10th coefficient reduce minimum temperature by 1 to 3 degrees Celsius, last two coefficients use continuous measures of rain and temperature for each hour between 8 am and 9 pm on the protest day respectively. Kleinbergen Paap F-Statistics reported right of each coefficient. Horizontal lines represent 90% and 95% confidence intervals.

Figure A.4. Baseline result with LASSO-selected instrumental variable



Note: Coefficient plot for 2SLS estimation of baseline regression using LASSO selected IV (stata command: `ivlasso`). IV with PDS-selected variables and full regressor set. High-dimensional instrument with 1,232 variables. These include rain and temperature for each hour between 8 am and 9pm as well as its interaction with itself, each other and the occurrence of a Monday protest. First coefficient from specification with no fixed effects and no controls. Second coefficient from specification with week FE, municipality and month of the year FE, as well as municipality linear time trends. Last coefficient from specification with all fixed effects and set of baseline controls. Selected instruments in model 1: rain 12pm \times protest, rain 1pm \times protest, temperature 2pm \times protest. Selected instruments in model 2: rain 12pm \times protest and temperature 2pm \times protest. Selected instruments in model 3: rain 11am \times protest, rain 1pm \times protest, temperature 2pm \times protest. Horizontal lines represent 90% and 95% confidence intervals.

Table A.1. Summary Statistics for Ever-Treated and Never-Treated Samples

	Ever Treated				Never Treated			
	Mean	SD	Min	Max	Mean	SD	Min	Max
Hate crime (probability)	0.0810	0.2728	0	1	0.0023	0.0482	0	1
Protest (probability)	0.0195	0.1383	0	1	—	—	—	—
Total participants	2.8729	32.2365	0	2300	—	—	—	—
Log(1+ participants)	0.0903	0.6498	0	7.7411	—	—	—	—
Total participants protest	147.2293	179.0992	20	2300	—	—	—	—
Pleasant weather × Protest	0.0157	0.1243	0	1	—	—	—	—
Pleasant weather (probability)	0.7564	0.4293	0	1	0.7550	0.4301	0	1
GDP per capita	3.13e+04	1.23e+04	1.92e+04	8.55e+04	3.03e+04	8.63e+03	1.51e+04	1.88e+05
Population density	715.5033	823.1041	41.6208	4736.1055	181.8333	273.0904	2.2124	4601.1714
Unemployment rate	0.0437	0.0149	0.0124	0.1003	0.0209	0.0129	0	0.2903
Share of refugees	0.0142	0.0091	0.0008	0.0493	0.0123	0.0060	0.0008	0.1300
AfD vote share	0.1413	0.0828	0.0352	0.3502	0.0993	0.0619	0.0222	0.3502
Share of total crime cases	0.6529	1.7095	0.0613	16.2552	11.2428	24.7601	0.0371	1182.4286
Municipality-week observations	21,678			2,764,730				

Table A.2. **First Stage: pleasant weather and protest participation**

	Log(1+ participants)			
	(1)	(2)	(3)	(4)
Pleasant weather \times Protest	0.317*** (0.0705)	0.252*** (0.0767)	0.247*** (0.0760)	0.250*** (0.0762)
Protest	4.336*** (0.429)	4.894*** (0.332)	5.022*** (0.364)	5.075*** (0.371)
Weather	-0.000699 (0.00607)	-0.000472 (0.00682)	-0.00175 (0.0107)	-0.000849 (0.0108)
Observations	21,678	21,678	21,678	21,678
Municipalities	84	84	84	84
Mean dep. var.	0.0903	0.0903	0.0903	0.0903
F first	20.26	10.76	10.54	10.77
Week FE	Yes	Yes	Yes	Yes
Municipality \times MoY FE	Yes	Yes	Yes	Yes
Municipality linear trends	Yes	Yes	Yes	Yes
Lagged controls	Yes	Yes	Yes	Yes
Economic controls	Yes	Yes	Yes	Yes
Refugee share		Yes	Yes	Yes
Right-wing share			Yes	Yes
Crime rate				Yes

Note: First stage estimation with week and municipality-month of year fixed effects and municipality linear time trends. Ever-treated sample only. SE clustered by municipality; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Time horizon is January 2015 until December 2019. Outcome is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Instrument is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleinbergen-Paap F-Statistics are reported at the bottom of the table. Controls are interacted with protest and weather. Controls in column 1 comprise GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, columns 2 to 5 add refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants respectively, all measured at the end of each calendar year.

Table A.3. Evidence against reporting bias

	State-week FE (1)	Region-week FE (2)	antisemitic HC (3)	any suspect (4)	suspect ratio (5)	arson ratio (6)	police employees (7)	police employees (8)
Log(participants)	0.536** (0.255)	0.501* (0.289)	-0.142 (0.109)	0.0321 (0.0413)	0.0115 (0.0177)	-0.0166 (0.0199)	-0.000545 (0.00164)	
Pleasant day protests								0.0165 (0.0536)
Observations	21316	19760	21678	21678	21678	21678	60	60
Unit Observations	82	76	84	84	84	84	12	12
Mean dep. var.	0.0875	0.0817	0.00932	0.00300	0.00171	0.00129	18.64	18.64
F first	11.68	9.445	10.77	10.77	10.77	10.77		
Municipality × MoY FE	Yes	Yes	Yes	Yes	Yes	Yes		
Municipality linear time trend	Yes	Yes	Yes	Yes	Yes	Yes		
All controls	Yes	Yes	Yes	Yes	Yes	Yes		
Week FE			Yes	Yes	Yes	Yes		
State × Week FE	Yes							
NUTS2 × Week FE		Yes						
Year FE							Yes	Yes
State FE							Yes	Yes

Note: Columns (1)–(6) report second-stage IV estimates where $\text{Log}(\text{participants})$ is instrumented by the number of protests on days with pleasant weather. $F \text{ first}$ reports the Kleibergen–Paap first-stage F-statistic. All specifications in columns (1)–(6) include municipality-by-month-of-year fixed effects, linear municipality time trends, and the additional fixed effects shown at the bottom. We include all baseline controls as well. Columns (7)–(8) show OLS results without an instrument, where the independent variables are summed over all municipalities in a given state and year. The dependent variables are indicated at the top of each column. Standard errors (in parentheses) are clustered at the municipality level for columns (1)–(6) and at the state level for columns (7)–(8). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4. **Football match types and hate crimes**

	any hate crime						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
any match	-0.0233 (0.0251)	-0.0228 (0.0249)	-0.0216 (0.0250)	-0.0237 (0.0244)	-0.0250 (0.0242)	-0.0243 (0.0259)	-0.0217 (0.0250)
log(attendance+1)	0.0031 (0.0028)	0.0030 (0.0028)	0.0029 (0.0028)	0.0033 (0.0027)	0.0035 (0.0027)	0.0032 (0.0029)	0.0029 (0.0028)
hooligan match				-0.0037 (0.0050)			
hooligan match certain					-0.0100 (0.0072)		
derby						-0.0101 (0.0126)	
contested match							-0.0126 (0.0299)
Observations	177284	177284	177284	177284	177284	177284	177284
Municipalities	94	94	94	94	94	94	94
Adj. R-squared	0.0829	0.0838	0.0883	0.0883	0.0883	0.0883	0.0883
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of the week × Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lagged matches		Yes	Yes	Yes	Yes	Yes	Yes
Lagged hate crimes		Yes	Yes	Yes	Yes	Yes	Yes
Municipality linear trends			Yes	Yes	Yes	Yes	Yes

Note: SEs clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until February 2020 (included). Unit of observation is municipality by day. Treatment is the occurrence of a football match. *had_hooligan_match* is a dummy equal to 1 if at least one team with a hooligan fan base playing a in a given municipality on that day. Hooligan definition is based on at least one source mentioning the team having hooligan fans. *had_hooligan_match_certain* is defined as at least two sourced mentioning the team. Outcome variable is equal to 1 if a hate crime took place in that municipality on the same day. *had_derby* is a dummy set to 1 when a match is considered a derby. *contested* is a dummy for matches that had extra time and/or ended with penalties. *has_hate_crimes* is a dummy for any hate crimes occurring on the same day as the match. All columns depict a two-way fixed effects regression with municipality by weekday and date fixed effects. Starting from column (2), we control for the count of matches and hate crimes on $t - 1$. From column (3) onwards, we allow for municipalities to be on linear daily trends. The mean of the dependent variable is 0.0190 with a standard deviation of 0.137. All specifications include 94 municipalities amounting to 177284 total observations.

Table A.5. **Football matches and hate crimes**

	any hate crime				number hate crimes	arson	rally	assault
	t	t+1	t+3	t+7	t+7	t	t	t
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
any match	-0.0216 (0.0250)	-0.0217 (0.0252)	0.0603* (0.0328)	0.0024 (0.0398)	0.0404 (0.0554)	-0.0011 (0.0028)	0.0013 (0.0056)	0.0017 (0.0090)
log(attendance+1)	0.0029 (0.0028)	0.0026 (0.0027)	-0.0062* (0.0033)	0.0000 (0.0042)	-0.0041 (0.0059)	0.0001 (0.0003)	-0.0002 (0.0006)	-0.0000 (0.0009)
Observations	177284	177284	177284	177284	177284	177284	177284	177284
Municipalities	94	94	94	94	94	94	94	94
Adj. R-squared	0.0883	0.138	0.193	0.250	0.350	0.00633	0.0355	0.0188
Mean dep. var.	0.0190	0.0356	0.0643	0.111	0.155	0.000287	0.000823	0.00356
SD dep. var.	0.137	0.185	0.245	0.314	0.659	0.0170	0.0287	0.0595
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of the week \times Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lagged matches	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lagged hate crimes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality linear trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: SEs clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until February 2020 (included). Unit of observation is municipality by day. Treatment is the occurrence of a football match on day t . *has.hate.crimes* is a dummy for any hate crimes occurring on the same day as the match in the given municipality. All columns depict a two-way fixed effects regression with municipality by weekday and date fixed effects. *has_hc_t1* is a dummy for any hate crimes occurring on the same day or the day after. *has_hc_t3* is a dummy for any hate crimes occurring up to 3 days after the match. *has_hc_t7* is a dummy including a 7-day window from match day. *sum.hate.crimes* is the count of hate crimes occurring within a 7-day window from match day. All columns also control for the count of matches and hate crimes on $t - 1$ as well as allowing for municipalities to be on linear daily trends. All specifications include 94 municipalities amounting to 177284 total observations.

Appendix B: Robustness Checks

To add:

- Update all robustness checks descriptions
- Control for lagged pleasant protest
- Focus on municipalities that have both unpleasant and pleasant protest
- Focus on municipalities with just one "unpleasant" protest day
- Estimate for East vs. West separately

Spillovers and spatial correlation. Our analysis is at the municipality level, which is the finest geographical level at which data is available to us. Since our baseline estimation focuses on the sample of ever-treated municipalities (and these tend to be geographically distant from each other) spill-overs across these locations are unlikely. However, there may be spatial spill-overs across municipalities, such that hate crimes in close by municipalities decreases while those in the protest location increase. While we address this concern directly in Table 2, we replicate our main analysis at higher levels of aggregation to account for those spillovers in Table B.1. We report our baseline estimates in column 1 and show OLS, 2SLS and Reduced Form estimates in Panels A, B, and C respectively. In columns 2 and 3 of Table B.1, we collapse our dataset to higher geographical administrative levels (NUTS-2 and NUTS-3 instead of districts for column 2 and 3 respectively). NUTS-3 regions correspond to cities and their sub-urbs in most cases. NUTS-2 regions are quite broad and capture entire states in some cases. Our results become slightly larger and magnitude and remain precisely estimated at the NUTS-3 level, but smaller and more noisily estimated for NUTS-2 regions, which is to be expected since there are only 18 regions and the F-stat falls below the conventional threshold. In a next step, we allow the spatial correlation of our observation within certain spatial windows using Conley standard errors (Conley, 1999). Columns 4 to 7 of Table B.1 show that our results become slightly noisier but are still significant when allowing correlation within 25km, 50km, 100km or 150km.

Additional controls. We verify in Table B.2 that our results hold when adding more controls. We successively include demographic controls, variables related to potential perpetrators and those related to social media. Note that in the OLS, these variables do not change the results because they do not vary much across the observation period and captured by the municipality fixed effects and linear time trends, but they may make a difference for the 2SLS and reduced form since we interact them with the weather and protest dummy variables as included instruments. Specifically, our demographic controls include: share of unskilled; share of females; dummies for share of population aged 0-25 25-50 and 50-75. Refugee controls include: share of foreign unemployed; share of skilled foreigners; share of asylum status granted. Social media controls include: the baseline Twitter penetration at the NUTS-3 level in 2014, the number of tweets containing #refugeesWelcome at the NUTS-3 level per capital in 2014 and the number of Facebook users following AfD before 2015 scaled by population. All of these variables are measured at the end of each calendar year. Throughout our results are similar in magnitude and largely precisely estimated. However, the first stage becomes slightly weaker for the saturated set of controls. Nonetheless, the stability and precision of our reduced form coefficients suggest that these variables do drive our results.

Variation on the fixed effect structure. Our main analysis include municipality by month of the year and week fixed effects, as well as municipality specific linear time trends. We show in Table B.3 that the results remain unchanged when we include different sets of fixed effects. Panel A corresponds to the TWFE identification and panel B to the 2SLS estimation and panel C reports the reduced form.

Column 1 uses only municipality and week fixed-effects, column 2 replaces the municipality fixed effects by fixed effects of the municipality and month of year pair. Column 3 adds per-municipality time trends but not municipality fixed effects. Column 4 includes municipality fixed effects and state specific linear time trends. Column 5 combines these with municipality month of year fixed effects, and column 6 with per-municipality trends. Throughout our result remain similar in magnitude and precisely estimated.

TWFE with staggered treatment. A growing body of literature highlights the limitations of two-way fixed effects (TWFE) difference-in-differences (DiD) estimators in settings with staggered treatment adoption and heterogeneous treatment effects (e.g., De Chaisemartin & D’Haultfoeuille (2022), Wooldridge (2021), Roth et al. (2022), Goodman-Bacon (2021)). When treatment effects vary over time or across units, standard TWFE estimators can assign negative weights to some comparisons, potentially leading to misleading or even incorrectly signed estimates. In particular, TWFE aggregates treatment effects using implicit weighting schemes that can emphasize comparisons between already-treated and later-treated units, violating the standard parallel trends assumption and introducing bias.

In our context, the presence of staggered right-wing protest activity across municipalities over time raises concerns about the validity of TWFE estimates. Specifically, if the effect of protest size on hate crimes varies across municipalities or evolves dynamically, the standard TWFE estimator may place disproportionate weight on inappropriate comparisons, distorting the estimated average treatment effect. To assess the extent of this issue, we implement the adjustment proposed by De Chaisemartin & d’Haultfoeuille (2020), which explicitly decomposes the TWFE estimate into group-by-time comparisons and calculates the weights assigned to each.

Table B.4 presents the results. Across all specifications, we find that the vast majority of the weight is positive, with 408 valid treatment-control comparisons contributing to the estimate. However, we detect a small fraction of negative weights—specifically, 15 instances, accounting for approximately -0.4% of the total sum of weights. This suggests that while TWFE does not produce extreme distortions in our case, there remains a minor degree of contamination from potentially invalid comparisons. Importantly, the stability of the estimates across specifications indicates that these negative weights do not significantly bias our main results. We also show later in a series of event-studies that our results hold when considering the first pleasant Monday PEGIDA protest only.

Nonlinearity, poisson and logit. Our main specification uses a linear probability model to estimate the effect of the log-transformed number of protest participants (plus one) on the occurrence of hate crime. We check that our results are robust to using non-linear models and to alternative definitions of our treatment variable. First, we define the treatment in three different ways: as the inverse hyperbolic sine transformation of participants in column 1, the number of participants as a share of the overall population of the municipality in column 2, and the absolute number of participants in column 3. Throughout our results suggest a positive and significant relationship between participation and hate crimes in the OLS and 2SLS estimation (the reduced form does not change, of course). Next, we estimate a poisson model to account for many zeros in the outcome variable. In fact, the likelihood of observing a hate crime in any municipality and week lies at about 8%; during protest weeks this is even higher. Nonetheless, we perform a Poisson pseudo-maximum likelihood regression in column 4 with the `ppmlhdfc` Stata command following Correia et al. (2019). This approach is not possible for the 2SLS estimation because of the incidental parameter problem. Consequently, we follow Lin & Wooldridge (2019) and implement a control function approach and robust standard errors in the second stage. We find in OLS, 2SLS and reduced form regressions that our results hold and are even more precisely estimated in the 2SLS. Next, column 5 estimates a logit model for the OLS specification, with no corresponding 2SLS estimation due to the same incidental parameter problem. Again, our results hold in all estimations.

Robustness to outliers. In Figure B.1, we examine the possibility that our results are driven by some municipalities that react very strongly to large protests and that in some weeks the salience of the refugee issue and resulting sensitivity to protest size is particularly high. In the top panel, we re-estimate our baseline specification and plot the 2LS coefficients and confidence intervals of our treatment variable dropping one municipality at a time. Our results are very similar in size and precision throughout, suggesting that no single municipality is driving the effect. The bottom panel shows the same estimates dropping single weeks, further confirming the robustness of our estimates.

First PEGIDA protest. Next, we examine whether the timing of the *first* PEGIDA protest is preceded by rising hate crime trajectories across municipalities. If the timing of the first protest is endogenous and has ripple effects into subsequent weeks, the event study may not capture visible pre-trends. In addition, this exercise alleviates concerns about multiple treatment of the same municipality. We show the event study results at the municipality-day and municipality-week level in Figure B.2. Consistent with our baseline event-study approach we find no differences in hate crime trajectories in the 7 days or 5 weeks leading up to the first Monday PEGIDA protest and a persistent increase in the level of hate crimes in subsequent days and weeks.

Controlling for weather on subsequent days. Crime rates are known to be influenced by weather (Chersich et al., 2019), and weather is temporally correlated. While the weather dummy should capture the average serial correlation of weather on hate crimes, one may imagine that subsequent weather impacts hate crimes differently when a protest occurred in the same week. To check that this does not influence our results, we turn to our event-study specification at the municipality and day level, and control for pleasant weather on each subsequent day. The left panel of Figure B.3 shows that our results remain unchanged when accounting for serial correlation in weather more explicitly.

Controlling for past cumulative protest. Our main specification controls for the occurrence of protest and hate crimes, as well as the number of participants in the previous week. In the event study, we do not control for lagged variables since the treatment effect is always relative to the level in $t-1$. However, it is possible that the treatment effect changes depending on whether the pleasant day protest happens later in the life-cycle of the PEGIDA movement or early on. Thus in the right panel of Figure B.3, we control for the cumulative number of protests (both Monday protest but also other PEGIDA protests) until $t-1$. Again, we find that our results remain similar in magnitude but marginally more noisily estimated.

Table B.1. **Robustness: Protest participation and hate crime - spatial correlation**

	Any hate crime in the same week						
	baseline (1)	nuts3 (2)	nuts2 (3)	Conley 25km (4)	Conley 50km (5)	Conley 100km (6)	Conley 150km (7)
Panel A: OLS							
Log(participants)	0.0192*** (0.00494)	0.0191*** (0.00505)	0.0174*** (0.00560)	0.0187*** (0.00472)	0.0187*** (0.00465)	0.0187*** (0.00468)	0.0187*** (0.00469)
Panel B: 2SLS							
Log(participants)	0.513** (0.215)	0.625** (0.261)	0.233 (0.222)	0.513* (0.295)	0.513* (0.289)	0.513* (0.290)	0.513* (0.298)
Panel C: Reduced Form							
Pleasant weather × Protest	0.128*** (0.0308)	0.138*** (0.0336)	0.0557 (0.0608)	0.128** (0.0559)	0.128** (0.0555)	0.128** (0.0548)	0.128** (0.0556)
Observations	21678	15958	4518	21678	21678	21678	21678
Number of geo. units	84	62	18	84	84	84	84
F first	10.77	10.19	7.114	10.77	10.77	10.77	10.77
Mean dep. var.	0.0810	0.110	0.389	0.0827	0.0827	0.0827	0.0827
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geo. unit × MoY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geo. unit linear time trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes
All controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: TWFE estimation with week and municipality-month of year fixed effects and municipality linear time trends. Ever-treated sample only. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Outcome comes from ProAsyl and AAF and is measured as a dummy variable equals to 1 if any hate crime against refugees was committed between Monday and Sunday of the same week. Instrument in Panel B is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleinbergen-Paap F-Statistics are reported at the bottom of the table. Controls comprise GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants, all measured at the end of each calendar year. Panel B and C additionally include the interaction between controls and the weather as well as the protest dummy.

Table B.2. **Protest Participation increases probability of hate crimes - more controls**

	Any hate crime in the same week			
	(1)	(2)	(3)	(4)
Panel A: OLS				
Log(participants)	0.0192*** (0.00498)	0.0184*** (0.00485)	0.0185*** (0.00485)	0.0185*** (0.00485)
Panel B: 2SLS				
Log(participants)	0.526** (0.227)	0.584** (0.267)	0.680** (0.327)	0.610 (0.652)
Panel C: Reduced Form				
Pleasant weather \times Protest	0.130*** (0.0316)	0.139*** (0.0345)	0.147*** (0.0338)	0.139*** (0.0347)
Observations	21678	21678	21678	21678
Municipalities	84	84	84	84
Mean dep. var.	0.0915	0.0915	0.0915	0.0915
F first	10.05	8.582	6.933	8.296
Week FE	Yes	Yes	Yes	Yes
Municipality \times MoY FE	Yes	Yes	Yes	Yes
Municipality linear trends	Yes	Yes	Yes	Yes
Lagged controls	Yes	Yes	Yes	Yes
Base controls	Yes	Yes	Yes	Yes
Demographic controls		Yes	Yes	Yes
Perpetrator controls			Yes	Yes
Social media controls				Yes

Note: Differences in differences estimation with week and municipality-month of year fixed effects. SE clustered by municipality; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Time horizon is January 2015 until December 2019. Treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Outcome comes from ProAsyl and AAF and is measured as a dummy variable equals to 1 if any hate crime against refugees was committed between Monday and Sunday of the same week. Instrument in Panel B is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleinbergen-Paap F-Statistics are reported at the bottom of the table. Municipality-level base controls include: the lagged log (1 +) number of protest participants in the previous week and the lagged dependent variable, GDP per capita; population density; unemployment share; refugee share, latest national vote shares for the AfD and crime rate per 100K inhabitants. Demographic controls include: share of unskilled; share of females; dummies for share of population aged 0-25 25-50 and 50-75. Refugee controls include: share of foreign unemployed; share of skilled foreigners; share of asylum status granted. Social media controls include: the baseline Twitter penetration at the NUTS-3 level in 2014, the number of tweets containing #refugeesWelcome at the NUTS-3 level per capital in 2014 and the number of Facebook users following AfD before 2015 scaled by population (taken from Müller & Schwarz, 2021) at the groups of municipality level. All measured at the end of each calendar year. Panel B and C additionally include the interaction between controls and the weather as well as the protest dummy.

Table B.3. **Robustness: varying the fixed effects structure**

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS						
Log(participants)	0.0208*** (0.00468)	0.0196*** (0.00486)	0.0201*** (0.00479)	0.0205*** (0.00472)	0.0197*** (0.00487)	0.0197*** (0.00488)
Panel B: 2SLS						
Log(participants)	0.512** (0.208)	0.542** (0.229)	0.496** (0.203)	0.478** (0.232)	0.543** (0.259)	0.475** (0.232)
Panel C: Reduced Form						
Pleasant weather × Protest	0.130*** (0.0295)	0.137*** (0.0328)	0.124*** (0.0289)	0.126*** (0.0404)	0.144*** (0.0444)	0.124*** (0.0399)
Protest	0.279* (0.146)	0.281* (0.147)	0.307** (0.138)	0.441*** (0.131)	0.415*** (0.134)	0.405*** (0.134)
Weather	0.00577 (0.0310)	-0.0137 (0.0335)	0.00146 (0.0303)	-0.0178 (0.0369)	-0.0149 (0.0348)	-0.0193 (0.0361)
Observations	21678	21678	21678	21316	21316	21316
Municipalities	84	84	84	82	82	82
Mean dep. var.	0.0915	0.0915	0.0915	0.0875	0.0875	0.0875
F first	10.846	10.797	10.552	11.459	11.606	11.419
All controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes			Yes		Yes
Week FE	Yes	Yes	Yes			
Municipality × MoY FE		Yes			Yes	
Municipality linear time trend			Yes			Yes
State linear time trend				Yes	Yes	Yes

Note: SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Panel A: two-way fixed effects regression with municipality and week fixed effects. Main treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Panel B: two-way fixed effects regression with municipality-month of the year and week fixed effects. Main treatment is the interaction between pleasant weather conditions during a Monday and the presence of protest in that week. In this panel, we additionally include weather and protest separately. We also include the interaction of all the controls with the both weather conditions and the presence of protest. For both panels; outcome comes from ProAsyl and AAF and is measured as a dummy variable equals to 1 if any hate crime against refugees was committed between Tuesday and Sunday of the same week. Municipality-level controls include: the lagged log (1 +) number of protest participants in the previous week and the lagged dependent variable; GDP per capita; population density; unemployment share; refugee share; latest national vote shares for the AfD and crime rate per 100K inhabitants.

Table B.4. Chaisemartin and D'Haultfoeuille (2020) TWFE negative weights

	Any hate crime in the same week							
	(1)		(2)		(3)		(4)	
OLS								
Log(participants)	0.0196*** (0.00500)		0.0194*** (0.00496)		0.0192*** (0.00493)		0.0192*** (0.00493)	
Week FE	Yes		Yes		Yes		Yes	
Municipality × MoY FE	Yes		Yes		Yes		Yes	
Municipality linear trends	Yes		Yes		Yes		Yes	
Lagged controls	Yes		Yes		Yes		Yes	
Economic controls	Yes		Yes		Yes		Yes	
Refugee share			Yes		Yes		Yes	
Right-wing share					Yes		Yes	
Crime rate							Yes	
	N. ATTs	Sum weights	N. ATTs	Sum weights	N. ATTs	Sum weights	N. ATTs	Sum weights
Positive weights	408	1.004	408	1.004	408	1.004	408	1.004
Negative weights	15	-.003974	15	-.00397	15	-.004127	15	-.00414

Note: Differences in differences estimation with week and municipality-month of year fixed effects and municipality linear time trends. Ever-treated sample only. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Outcome comes from ProAsyl and AAF and is measured as a dummy variable equals to 1 if any hate crime against refugees was committed between Monday and Sunday of the same week. Controls in column 1 comprise GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, columns 2 to 5 add refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants respectively, all measured at the end of each calendar year.

Table B.5. **Protest Participation increases probability of hate crimes (without Monday)**

	Any hate crime in the same week			
	(1)	(2)	(3)	(4)
Panel A: OLS				
Log(participants)	0.00849** (0.00375)	0.00826** (0.00370)	0.00810** (0.00368)	0.00808** (0.00368)
Panel B: 2SLS				
Log(participants)	0.282** (0.136)	0.361* (0.201)	0.377* (0.205)	0.377* (0.201)
Panel C: Reduced Form				
Pleasant weather × Protest	0.0899** (0.0342)	0.0910** (0.0355)	0.0932*** (0.0348)	0.0945*** (0.0345)
Protest	0.233** (0.0943)	0.217* (0.112)	0.165 (0.106)	0.183* (0.108)
Weather	-0.0168 (0.0216)	-0.0159 (0.0218)	-0.0177 (0.0332)	-0.0180 (0.0343)
Observations	21,678	21,678	21,678	21,678
Municipalities	84	84	84	84
Mean dep. var.	0.0810	0.0810	0.0810	0.0810
F first	19.97	10.67	10.46	10.69
Week FE	Yes	Yes	Yes	Yes
Municipality × MoY FE	Yes	Yes	Yes	Yes
Municipality linear trends	Yes	Yes	Yes	Yes
Lagged controls	Yes	Yes	Yes	Yes
Economic controls	Yes	Yes	Yes	Yes
Refugee share		Yes	Yes	Yes
Right-wing share			Yes	Yes
Crime rate				Yes


Note: Differences in differences estimation with week and municipality-month of year fixed effects and municipality linear time trends. Ever-treated sample only. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Treatment is measured as the log of 1 + total participants at a PEGIDA Monday demonstration using data from Kanol & Knoesel (2021). Outcome comes from ProAsyl and AAF and is measured as a dummy variable equals to 1 if any hate crime against refugees was committed between Tuesday and Sunday of the same week. Instrument in Panel B is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleinbergen-Paap F-Statistics are reported at the bottom of the table. Controls in column 1 comprise GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, columns 2 to 5 add refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants respectively, all measured at the end of each calendar year. Panel B and C additionally include the interaction between controls and the weather as well as the protest dummy.

Table B.6. **Robustness: treatment definition and non-linear estimation**


	Any hate crime in the same week				
	asinh (1)	ratio (2)	absolute (3)	poisson (4)	logit (5)
Panel A: OLS					
IHS(participants)	0.0169*** (0.00434)				
Participants (as share of pop.)		3.850* (2.022)			
Participants (absolute number)			0.000354*** (7.92e-05)		
Log(participants)				0.0825** (0.0379)	0.168*** (0.0585)
Panel B: 2SLS					
IHS(participants)	0.508** (0.213)				
Participants (as share of pop.)		85.16* (46.23)			
Participants (absolute number)			0.00375** (0.00183)		
Log(participants)				0.0872*** (0.0273)	
Panel C: Reduced Form					
Pleasant weather × Protest	0.128*** (0.0308)	0.128*** (0.0308)	0.128*** (0.0308)	0.677*** (0.204)	1.065** (0.473)
Protest	0.317** (0.139)	0.317** (0.139)	0.317** (0.139)	1.170 (1.604)	3.934* (2.046)
Weather	-0.0193 (0.0328)	-0.0193 (0.0328)	-0.0193 (0.0328)	-0.175 (0.499)	-0.348 (0.625)
Observations	21678	21678	21678	13838	13838
Municipalities	84	84	84	80	80
Adj. R-squared	0.196	0.196	0.196	.	.
Mean dep. var.	0.0915	0.0915	0.0915	0.143	0.143
Week FE	Yes	Yes	Yes	Yes	Yes
Municipality × MoY FE	Yes	Yes	Yes	Yes	Yes
Municipality linear time trend	Yes	Yes	Yes	Yes	Yes
All controls	Yes	Yes	Yes	Yes	Yes

Note: Differences in differences estimation with week and municipality-month of year fixed effects. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Outcome comes from ProAsyl and AAF and is measured as a dummy variable equals to 1 if any hate crime against refugees was committed between Monday and Sunday of the same week. Instrument in Panel B is pleasant weather dummy as defined in section 3.2 interacted with scheduled Monday protest, controlling for protest and weather separately. Kleibergen-Paap F-Statistics are reported at the bottom of the table. Controls include: the lagged log (1 +) number of protest participants in the previous week and the lagged dependent variable, GDP per capita; population density; unemployment share; refugee share, latest national vote shares for the AfD and crime rate per 100K inhabitants. Panel B and C additionally include the interaction between controls and the weather as well as the protest dummy.

Figure B.1. Robustness to excluding potential outliers




(a) Dropping municipalities



(b) Dropping weeks

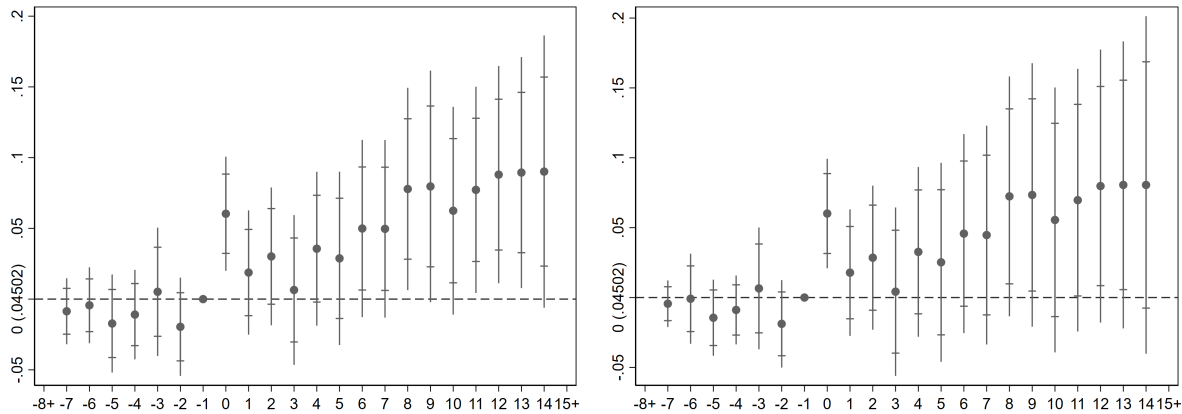
Note: Coefficient plot for 2SLS estimation of baseline regression for different sample compositions. Top panel shows point estimates (vertical lines represent 90% and 95% confidence intervals) when dropping one municipality at a time. Bottom panel does the same but dropping one week at a time. There are a total of 84 ever treated municipalities and 260 weeks.

Figure B.2. Event study: first PEGIDA protest and hate crimes at the day-level (left panel) and week-level (right panel)



Note: Event study using `xtevent` from Freyaldenhoven et al. (2024). Regression at the municipality and day level with linear time trends where outcome is dummy for any hate-crime on day t (left panel) or week t (right panel). Estimating equation: $\text{hate crime}_{it} = \sum_{k=T_0}^{-1} \beta_k \text{pleasant} \times \text{protest}_{ik} + \sum_{k=0}^{T_1} \beta_k \text{pleasant} \times \text{protest}_{ik} + \eta_1 \text{protest}_{it} \times X'_{it} + \eta_2 \text{weather}_{it} \times X'_{it} + \gamma_1 X'_{it} + \gamma_2 \text{protest}_{it} + \gamma_3 \text{pleasant}_{it} + \mu_{im} + \delta_t + \epsilon_{it}$. Standard errors are clustered at the municipality level. Set of controls X_{it} includes all controls of Table 1 column 4 except for the lagged controls, as well as dummy for post protest and post pleasant weather period. Sample consists of ever-treated municipalities only. Vertical bars present 90% and 95% confidence intervals, respectively. Left panel includes rain and temperature between 8 am and 9 pm on each day. Right panel includes the cumulative number of protests in the municipality until previous day.

Figure B.3. Event study: PEGIDA protest and hate crimes at the day-level controlling for weather on each day (left panel) and lagged cumulative protest (right panel)



Note: Event study using `xtevent` from Freyaldenhoven et al. (2024). Regression at the municipality and day level with linear time trends where outcome is dummy for any hate-crime on day t . Estimating equation: $\text{hate crime}_{it} = \sum_{k=T_0}^{-1} \beta_k \text{pleasant} \times \text{protest}_{ik} + \sum_{k=0}^{T_1} \beta_k \text{pleasant} \times \text{protest}_{ik} + \eta_1 \text{protest}_{it} \times X'_{it} + \eta_2 \text{weather}_{it} \times X'_{it} + \gamma_1 X'_{it} + \gamma_2 \text{protest}_{it} + \gamma_3 \text{pleasant}_{it} + \mu_{im} + \delta_t + \epsilon_{it}$. Standard errors are clustered at the municipality level. Set of controls X_{it} includes all controls of Table 1 column 4 except for the lagged controls, as well as dummy for post protest and post pleasant weather period. Sample consists of ever-treated municipalities only. Vertical bars present 90% and 95% confidence intervals, respectively. Left panel includes rain and temperature between 8 am and 9 pm on each day. Right panel includes the cumulative number of protests in the municipality until previous day.

Appendix C: Data Appendix

To add:

- Complete list of variable sources
- More details on Twitter sentiment analysis

C.1 PEGIDA Protest Data

To create this dataset, the Kanol & Knoesel (2021) identified relevant parliamentary questions that contained information on right-wing extremist demonstrations. They then extracted the relevant data from tables included in these responses and merged them to create a comprehensive dataset. To classify each demonstration based on its ideology, an identification variable was added to the dataset. This classification process was based on descriptions provided in the government’s responses to parliamentary questions. Demonstrations were classified as ”right-wing extremist,” ”mostly right-wing,” or ”partially right-wing” based on these descriptions. The number of right-wing protests was highest in 2015 (with 290 demonstrations) and lowest in 2010 (with only 70 demonstrations). Of all demonstrations in this dataset, over 83% were classified as ”right-wing extremist,” while around 17% were categorized as ”mostly right-wing.” Only a very small fraction (0.2%) was identified as ”partially right-wing.”

The authors also used geocoding techniques to identify the location of each demonstration. This involved converting textual descriptions of locations into geographic coordinates that could be plotted on a map. Some demonstrations were held in more than one place or moved through multiple locations. We treat these protests as separate incidents. In some cases, exact numbers for participants in a demonstration were not available; instead, an estimation was given (e.g., 5-10 or 100-500). In these cases, we follow the authors and use the average of this range of numbers.

C.2 Spatial Harmonization

Our regional-level of analysis is the municipality level as defined above. However some variables are available only at the district level, the values of these variables remain identical across municipalities within the same district. Both municipalities’ and districts’ borders changed during our sample period, i.e., 2015 - 2020. Hence, for reasons of data coherence, we adjusted all variables according to one border division.

Municipality-level variables are adjusted to the border division as of December 2020. We use the ’name and area changes of municipalities’ tables which are published yearly by DESTATIS, which document four types of municipality changes: (1) municipalities that merged with other municipalities, or joined an existing municipality; (2) municipalities that split to several municipalities; (3) change of key; (4) change of name. Municipality-level variables from before December 2020 were updated as follows. First, names and keys were updated to December 2020. Second, for merged municipalities (i.e., change (1)), averaged variables (e.g., voting turnout) were updated as a population-weighted average, for municipality i in year $t=2015, \dots, 2020$, and the new merged municipality j : $Var_{j,t} = \sum_{i=1, \dots, n} Var_{i,t} * \frac{Pop_{i,t-1}}{Pop_{j,t}}$; and for summed variables (e.g., total votes for AfD) $Var_{j,t} = \sum_{i=1, \dots, n} Var_{i,t}$. The split municipalities (i.e., change (2)) were dropped due to their small share in the sample, with less than half a percent of all municipalities.

District-level variables are adjusted to the NUTS3 2013 version, which entered into force on 31 December 2013 and applied from 1 January 2015. Two changes were made since then in 2017: (1) The border between Cochem-Zell and Rhein-Hunsrück-Kreis slightly shifted, without affecting other districts borders; (2) Göttingen and Osterode am Harz merged into one district under the name of Göttingen. Change (1) was ignored, since the boundary shift is minor in terms of km^2 area. To account

for change (2), all regional controls in Göttingen and Osterode am Harz after 2016, received a value equal to the value in Göttingen, weighted by the share of the corresponding district in 2016, s.t.: for averaged variables (e.g., unemployment rate) $Var_{i,t} = Var_{i,t} * \frac{Var_{i,2016}}{Var_{j,2016}}$; and for summed variables (e.g., population) $Var_{i,t} = Var_{i,t} * \frac{Var_{i,2016}}{Var_{i,2016} + Var_{j,2016}}$, with $i, j \in [\text{Göttingen}, \text{Osterode am Harz}]$ and $t \in [2017, 2020]$ Moreover, GDP data was available only in the NUTS3 2016 format (i.e., also before 2017). Hence only GDP was recovered weighted by the population share of the two regions.

C.3 Weather Data

We take information on weather conditions on protest days from ERA5, which is a global atmospheric reanalysis dataset produced by the European Centre for Medium-Range Weather Forecasts (ECMWF). ERA5 provides a rich historical record of global weather conditions dating back to 1979 from multiple sources, including satellites, radiosondes, and weather stations. ERA5 includes hourly information on a variety of meteorological variables, including temperature, humidity, wind speed, precipitation, and atmospheric pressure, among others. These data are presented at a high spatial and temporal resolution. We extract information on precipitation (rain in mm) and temperature for every hour during protests times, i.e. on Mondays between 12pm and 5pm to create our indicator for pleasant weather.

C.4 LASSO IV

Our context allows us to address many challenges associated with exploiting local weather shocks to instrument for protest. However, we still face the challenge of multi-dimensionality: we must consider a wide range of potentially relevant and interacting weather conditions across various municipalities. Weather conditions might influence protest participation through multiple channels. Fundamentally, adverse weather conditions may deter protest attendance due to increased discomfort or logistical challenges. Conversely, favorable weather conditions might facilitate higher turnout by reducing the costs of participation (Madestam et al., 2013). In addition, favorable weather conditions may draw in a higher number of media outlets or generate a more positive coverage (Zhong & Zhou, 2012) and influence how much participants associate the event with positive sentiments (Goetzmann et al., 2015; Jiang et al., 2022).

In order to address this complexity and decrease the researchers' degree of freedom as well as the risk of finding false positives, we implement a LASSO IV approach similar to Beraja et al. (2023). We take hourly information on precipitation and temperature between 8 am and 9 pm on the protest day and allow each variable to interact with each other and with the protest dummy. We implement a LASSO regression that selects predictors of protest participation based on over 1200 variables. Standard errors are calculated using the cross-fit partialing-out LASSO IV algorithm following Chernozhukov et al. (2018). Appendix Table ?? reports the LASSO selected instruments in the first stage regression.

C.5 Social Media Data

Overall Twitter Use To estimate the overall Twitter penetration at baseline in Germany, we used the Twitter Academic Research API to sample tweets containing identified by Twitter as being written in German, and containing at least one of the 100 most frequent words in German.²³ For each of these tweets, we used the user's stated location (if available). The location field is part of the user's profile and can be filled with any text the user wants. We only get the location indicated at the time of collection (December 2020). We geolocalize these text locations using Open Street Map's Nominatim geocoder to a NUTS-3 region, and discard any location that is too vague (corresponding to a larger region or

²³Since the Twitter API does not allow to search directly for all tweets in German, we search for tweets containing the 100 most frequent words in German, as listed by Sharoff (2006) on the website <http://corpus.leeds.ac.uk/frqc/>.

Germany as a whole), as well as locations outside Germany. Our sampling strategy is as follows: for each 3-hour interval in the years 2013 and 2014, we choose a random second in this interval, and ask for the first 100 tweets posted during this second. We obtained 577 000 tweets in this manner. For each request, the time difference between the requested instant and the last tweet returned allows estimating the rate of tweets posted at this moment, and from this we deduce an estimate of the rate of tweets in each NUTS-3 region. We then average these estimated rates over the 2 year period to obtain the overall Twitter use estimate.

Using a larger sample of tweets (up to end of 2018) collected in the same way,²⁴ we estimate the intensity of social media connections between NUTS-3 regions: we measure the social media influence of a region i on a region j by measuring the rate of tweets originally posted by users located in i retweeted by users in j .

PEGIDA tweets We collected all tweets in German and English containing the word PEGIDA posted between October 2014 and 2021. This dataset contains 2,068,258 tweets. We geolocalized these tweets using the same method as for overall Twitter use and obtained 659,709 geolocalized tweets and retweets along with their date of posting, and the original tweet in case of a retweet. We use these tweets to measure the social media influence between pairs of regions, but focusing on the PEGIDA-related network.

Pro-Refugee tweets To measure pro-refugee sentiment, we collected all tweets in German and English containing the hashtag #RefugeesWelcome and posted between 2013 and 2018, representing 390,000 tweets. We are able to geo-localize about 150,000 of these tweets to NUTS-3 regions in Germany.

Twitter Sentiment

C.6 GENIOS Newspaper Data

Newspaper articles talking about PEGIDA protests: We use the GENIOS database, which contains newspaper articles from 282 different publications between 2000 and 2023. We filter the pool of articles and select all articles that contain the word PEGIDA published between 2015 and 2019, then use a large language model (gemini-1.5-flash-002) with a custom prompt to determine whether the articles mention pro-PEGIDA protests or counterprotests. We build a dataset of these articles, along with the types of demonstrations mentioned, their publication date, and the newspaper that published them.

The detailed prompt and output format specification we give is the following, where {ArticleTitle}, {ArticleText}, {PublicationDate} and {PublicationName} are replace with the relevant values:

```
PROMPT = """
START OF ARTICLE
{ArticleTitle}

{ArticleText}
END OF ARTICLE

The article was published on {PublicationDate} in {PublicationName}.

Please list the PEGIDA (or PEGIDA-related) protests and counterprotests mentioned in
the article. Return an array of dictionaries, with one entry per demonstration. The
dictionary is structured as follows: {"location": "<city of the protest>",
"demonstration_type": "<PEGIDA or anti-PEGIDA>", "date": "YYYY/MM/DD",
"participants": <number of participants, or range like 2000-5000, or "unavailable">}.
Please include all fields. If there are no protests explicitly mentioned in the article
```

²⁴Due to the restriction on the Twitter API placed in March 2023, we are unable to complete this dataset to the end of 2020.

```

(or if it is very vague), simply answer with an empty array.
"""

FORMAT_DEFINITION = {
  "type_": "ARRAY",
  "items": {
    "type_": "OBJECT",
    "properties": {
      "location": {"type_": "STRING"},
      "demonstration_type": {"type_": "STRING"},
      "date": {"type_": "STRING"},
      "participants": {"type_": "STRING"}
    },
    "required": ["location", "demonstration_type", "date", "participants"]
  }
}

```

Newspaper diffusion data: To measure the diffusion of articles about PEGIDA protests and counter-protests, we use the 2020 IVW newspaper diffusion data. The diffusion of newspaper (or more precisely advertisement placement units²⁵) has been measured at the municipality level through surveying sales channels during one reference week in November 2019. Newspaper sales were then proportionally adjusted to match the sales number of the first quarter of 2020. Thus, we have the estimated diffusion for each ad placement unit and municipality. We match this data to the GENIOS data by manually matching the 893 advertising units to newspapers in the GENIOS database, when available.²⁶ We end up matching 485 advertising units, representing 84% of the total measured diffusion, to 134 publications reference in GENIOS. We obtain a dataset indicating, for each matched GENIOS publication, its diffusion in each municipality.

C.7 ChatGPT hate crime classification

To better understand who is committing hate crimes, what motivates them, whether the victims belong to particular demographic groups, and the type and location of hate crimes, we use ChatGPT to analyse hate crime descriptions. ChatGPT is a generative large language model (LLM), capable of efficient and swift advanced textual analysis. By fine-tuning the baseline model²⁷, we train the LLM to carefully dissect all available information on each hate crime and output a battery of descriptive variables.

The final prompt used for fine-tuning the responses and retrieving relevant information appears below. We first instruct the model to carefully analyse the text, followed by detailed instructions including examples of the expected output columns. The prompt returns 38 columns in a Python data frame for each hate crime. The precise text of the prompt is the result of careful prompt engineering efforts which balance the level of detail in output information against not identifying sufficient information. Overly specific prompts lead to the model not retrieving information on any columns; whereas excessive vagueness does not allow the model to understand what we are looking for.

```

def hatecrime_prompt_gpt4(text):
    input = [
        {
            "role": "system",

```

²⁵The IVW data is aimed at buyers of advertisement placement. Newspapers commonly offer multiple local editions, allowing buyers of ads to target their ads to finer-grained regions, rather than publishing their ad for the whole readership of the newspaper.

²⁶The GENIOS data does not indicate if an article was only published in a local edition of a paper. We assume here that papers with the same name will publish approximately the same articles, and that the main thing that changes are the advertisements presented, as this allows newspaper to get more money from advertisers by distributing ads that are more targeted and thus more effective.

²⁷The prompts were implemented using the version GPT-4O-MINI-2024-07-18. Using the flagship ChatGPT version at the time of this writing (GPT-4O-2024-08-06 implies significantly higher usage costs with no improvement in text classification).

```

"content": """
You are a helpful assistant that carefully reads through descriptions of
hate crimes and extracts as much information as possible.
Your only output is a short answer.
Act as an economist analysing crimes.
You are given a description of a crime.
Your task is to analyse the description and categorise the information
based
on several criteria that you are provided.
Manually fill in a template for each crime description.
First, work out how you would fill in each category.
Then evaluate your output.
If you are sure of it, paste it into the template.
Use the following step-by-step instructions:

1. Read the description.
2. Fill in the template.
3. Output a Python dataframe.
"""
},
{
"role": "user",
"content": text
},
{
"role": "user",
"content": """
I will now give you a list of categories.
Your first task is to evaluate whether this information is available for
the
given text. If yes, extract it.

1. **Bystanders_Present:** Note whether bystanders were mentioned in the
description
(e.g., "Yes", "No").
2. **Crime_Type:** Categorise the type of crime
(e.g., "Physical assault", "Verbal assault", "Arson").
3. **Criminal_History:** Indicate if the perpetrator's criminal history
is mentioned
(e.g., "None", "Minor offences", "History of hate crimes").
4. **Degree_of_Violence:** Describe the severity of the violence
(e.g., "Moderate", "Severe", "Minor").
5. **Immediate_Provocation:** Note any immediate provocations mentioned
in the description
(e.g., "Verbal altercation").
6. **Location_Type:** Describe the type of location
(e.g., "Train station", "Public street", "Private property").
7. **Motivation:** Extract the motivation behind the crime,
such as "Racism", "Xenophobia".
8. **Organisation_Level:** Indicate if the perpetrator was part of an
organised group or acted
spontaneously
(e.g., "Organised", "Spontaneous").
9. **Perpetrator_Anonymity:** Indicate if the perpetrator's identity is
known or unknown
(e.g., "Known", "Unknown").
10. **Perpetrator_Ethnicity:** Identify the perpetrator's ethnicity if
mentioned
(e.g., "Arab", "Caucasian").
11. **Perpetrator_Nationality:** Identify the perpetrator's nationality

```

- (e.g., "German").
12. ****Perpetrator_Religion:**** Identify the perpetrator's religion if mentioned (e.g., "Muslim", "Christian").
13. ****Perpetrator_Type:**** Describe the perpetrator (e.g., "Unidentified male").
14. ****Perpetrator_Age:**** List the age of the perpetrator if available (e.g., "30").
15. ****Perpetrator_Behaviour:**** Describe the behaviour of the perpetrator during the crime (e.g., "Aggressive", "Calm").
16. ****Perpetrator_Gender:**** Identify the gender of the perpetrator (e.g., "Male", "Female").
17. ****Perpetrator_Group:**** State if multiple perpetrators were involved or if the perpetrator acted alone (e.g., "Single", "Multiple").
18. ****Perpetrator_Motivation:**** Provide the likely motivation (e.g., "Racism", "Xenophobia").
19. ****Perpetrator_Political_Affiliation:**** Note any political affiliations if mentioned (e.g., "Far-right", "Neo-Nazi").
20. ****Public_Space:**** Indicate if the crime occurred in a public space ("TRUE" or "FALSE").
21. ****Radicalisation:**** Indicate if there is a possibility of radicalisation ("Possible", "Confirmed", or "NA").
22. ****Social Media Involvement:**** Indicate if social media played a role in the incident (e.g., "Yes", "No").
23. ****Specific Crime Type:**** Specify the type of hate crime (e.g., "Arson", "Verbal assault").
24. ****Symbols_Used:**** Note any symbols used during the crime (e.g., "Swastika").
25. ****Time_of_Day:**** Mention the time of day if available (e.g., "Morning", "Evening").
26. ****Victim_Ethnicity:**** Identify the victim's ethnicity if mentioned (e.g., "African", "Asian").
27. ****Victim_Religion:**** Identify the victim's religion if mentioned (e.g., "Muslim", "Christian").
28. ****Victim_Age:**** List the ages of all victims mentioned (e.g., "25", "30").
29. ****Victim_Gender:**** Identify the genders of all victims mentioned (e.g., "Male", "Female").
30. ****Victim_Nationality:**** Identify the nationalities of all victims mentioned (e.g., "Syrian").
31. ****Violence_Degree:**** Describe the degree of violence involved in the crime (e.g., "Severe", "Moderate").
32. ****Weapons_Used:**** Note any weapons used during the crime (e.g., "Knife", "Baseball bat").
33. ****Relationship:**** Indicate if the victim and perpetrator knew each other and how (e.g., "Strangers", "Acquaintances").
34. ****Previous Arrests or Convictions:**** Note any mention of the perpetrator's past legal encounters or arrests.

- 35. ****Extremist Group Affiliation:**** Indicate if the perpetrator is associated with any known extremist or neo-Nazi group.
- 36. ****Escalating Violent Behaviour:**** Identify if the description suggests a pattern of increasingly violent or extreme actions by the perpetrator.
- 37. ****Known by Law Enforcement:**** Indicate if the perpetrator was previously known to the police or law enforcement.
- 38. ****Past Involvement in Hate Movements:**** Describe if the perpetrator has a history of participating in hate-fuelled demonstrations or events.

Now fill in the template based on the text description of a hate crime below.

```
"""
},
{
    "role": "user",
    "content": text
},
{
    "role": "user",
    "content": """
```

Based on the text description of a hate crime, fill in the template for the categories outlined above.

After you fill in the template, provide the output in the form of Python code that creates a Pandas DataFrame (hc_df). Follow these guidelines strictly:

1. **Format Requirements:**

- All output values must be inside a list (e.g., 'Radicalisation': [[Possible]]).
- For any unknown or not specified information, use [[None]].
- For boolean responses, use [[True]] or [[False]].
- For numeric data, output the numbers directly in a list (e.g., [[50, 28, 30]]).

2. ****String Formatting**:**

- All strings should be parsed correctly and enclosed within single quotes.
- If the value is a boolean, None, or an integer: return it as is. If it is not, enclose it as string.
- Ensure no internal quotes in strings are present. Use appropriate replacements if needed. For example, 'Asylum seekers\' accommodation' should be written as 'Asylum seeker accommodation'.

3. ****Output Example**:**

- Your output should start with: hc_df =.
- The DataFrame creation should strictly follow this format:

```
hc_df = pd.DataFrame({
    'Bystanders_Present': [['No']],
    'Crime_Type': [['Physical assault']],
    ...
```

```

    })

    4. Important Rules:
    - Do not include any text outside of the DataFrame creation code.
    - The output must exactly match the provided text description.
    - Do not import any packages or include extra comments.

    Ensure that the output begins with hc_df = and follows the exact format
    guidelines provided above.

    """
    }
]

# Generate a response
response = client.chat.completions.create(
    model="gpt-4o-mini-2024-07-18",
    messages=input,
    temperature=0.2
)
output = response.choices[0].message.content.strip()
return output

```

The following example illustrates how the fine-tuned model responds to the previously outlined prompt. The automatised procedure takes as input a single hate crime description such as:

Um die Mittagszeit ist ein 27-jähriger Mann aus Somalia mit seinem Fahrrad unterwegs. Ein 48-jähriger bereits polizeibekannter Deutscher ruft ihm beleidigende Äußerungen zu und bewirft sein Fahrrad mit einer Fahrradkette und trifft ihn am Bein. Im Zuge der Ermittlungen stellt sich heraus, dass der Deutsche den Betroffenen bereits mehrfach beleidigt und beschimpft hatte.²⁸

For this particular hate crime, the LLM captures detailed information. The model records the crime as *"physical assault"* with a *"moderate"* degree of violence with a *"racist"* motivation. The location of the crime is reported as *"public street"*. The fine-tuned model also recognises several characteristics about the perpetrator, such as their gender (*"male"*), nationality (*"German"*), and age (*48*). Details about the perpetrator's past behaviour and radicalisation are also captured. Thus, the LLM recognises that the perpetrator has a criminal history and is known by law enforcement. Interestingly, the model classifies the perpetrator and victim as strangers, despite the text mentioning previous altercations between the same individuals.

C.8 Football Matches

To distinguish the impact of protest participation as a public signal of preferences for anti-immigrant policies from its role as a coordination device for far-right supporters, we use football matches as a quasi placebo.

Football matches, which naturally gather large groups without any explicit public signal of a political protest, offer an ideal setting to examine coordination among (right-wing) individuals in the absence of overt signalling. By comparing hate crime occurrences following football matches to the aftermath of PEGIDA protests, we aim to disentangle public signalling from mere coordination.

²⁸English translation: At lunchtime, a 27-year-old man from Somalia is travelling on his bicycle. A 48-year-old German, who is already known to the police, shouts insults at him and throws a bicycle chain at his bike, hitting him on the leg. In the course of the investigation, it emerged that the German had already insulted and abused the victim several times.

We further home in on the possible coordination occurring during football matches with right-wing fandom. If public signalling is the mechanism driving the variation in hate crimes, coordination between right-wing fans during football matches ought to yield negligible effects on hate crimes.

This approach enables us to robustly attribute changes in radical action to the public signal provided by successful PEGIDA protests, rather than to the mere opportunity for coordination that any large gathering (of like-minded individuals) might present.

Football matches. Football matches can be considered events where coordination among (right-wing) individuals can occur without an explicit public signal of their political preferences. To this end, we collected data on football match outcomes from the website *fbref.com*, which offers comprehensive statistics for various domestic and international football leagues. For Germany, data is available for three domestic male leagues as well as the national cup (*Bundesliga*, *2. Bundesliga*, *3. Liga*, and *DFB-Pokal*, respectively). Historical data availability varies across the competitions. The longest available time series includes all football matches from the 1988 season onwards for the first male league. Data for the domestic male cup is available starting from the 2014 season. The detailed data on football matches is collected via web scraping. All seasons spanning from 2014 to the season ending in 2024 are included in the initial dataset. The web scraping process starts by accessing the unique URLs of each league and each pertinent season. Each seasonal league page contains a section entitled "Scores & Fixtures" which includes a detailed table of all matches occurring in a given season, as well as their outcomes. We then extract the contents of these tables from the webpages. The extracted data includes match date and time, participating teams, scores, attendance, venue name, and additional unstructured notes such as whether the game ended with extra time or penalties.

Venue locations. In the next step, we use an open-source geocoding software from *OpenStreetMap*²⁹ to find the precise locations of each distinct venue. Mapping merely based on venue name, as provided in the *fbref.com* tables, is imprecise for several reasons. First, venues change names frequently. Venue names within *fbref.com* are harmonised over time. However, *OpenStreetMap* may continue using the previous name version or, on the contrary, a name change may not be reflected on *fbref.com*. Second, several venues in Germany are homonyms. For example, there are two Parkstadion-s which hosted a cup match. One is located in the state of *Hessen* and the other in *Mecklenburg-Vorpommern*. To distinguish name duplicates, we associate each venue with the name of the team that played there most frequently as the home team. Third, venue names may include with additional unstructured string descriptors which render them unrecognisable for a geolocating software. The venue *Audi-Sportpark*, for one, hosted 18 matches on artificial grass. The venue for these matches is named *Audi-Sportpark - Platz 5 (Kunstrasen)* which is unidentifiable for a geolocating software.

All venues that appear as the home venue for multiple teams (indicating potential name duplicates) or that cannot be geolocated due to possible name changes are flagged for manual review. We use *DFB.de* - German Football Association's official website - to search for the match information for matches whose venue cannot be identified. We then complement this data with a Wikipedia search to get the precise address of each venue. For venues that are still not geolocated, the script fills in missing geolocation details by using information from other entries associated with the same home team. This step solves the issue with *Audi-Sportpark* described above. Thereafter, we attempt to geolocate the venues using the precise address, as opposed to venue name. We successfully geolocate all venues following this multi-step iterative process.

Before the final step, we leverage the API offered by the open source tool *OpenCageData* to map venue coordinates to relevant NUTS regions and municipalities. To conclude, we merge each football match to venue geolocation data based on the venue name and home team³⁰. For any unmatched observations, we perform a second merge based on venue name alone. This step captures the cases where

²⁹We use the API provided by Nominatim which can be accessed directly through their website

³⁰Home team is defined as the team that had the highest number of matches on a given venue.

a team played on a venue that is not considered their home venue. A particularly prominent case is the final of DFB-Pokal which is traditionally always played on the Olympic Stadium of Berlin, independent of the participating teams.

Hooliganism and right-wing fanbase. To assess whether the variation in hate crimes as a result of PEGIDA protests is driven by the public signalling channel, we want to test the role of an occasion of pure coordination - a football match. For this purpose, we home in on the potential coordination occurring during football matches where at least one team has a fan base that can be categorised as far-right hooligan. We collect data on the potential hooligan tendencies of the fandom from a plethora of sources. Our primary source is Duben (2015) who highlights 9 clubs as having far-right fanbases. We complement that data with a report from the Federal Agency for Civic Education or bpb (acronym for the German name Bundeszentrale für politische Bildung) (Claus, 2024) and with a recent New York Times article by Hughes (2024). We verify the classification against the teams mentioned in relevant forum discussions on the social media platform Reddit.com.

Derby matches. Derbies are high-rivalry matches, generally occurring between two teams of the same city or between the top clubs of a country. The primary source for information on German derbies is Bundesliga. We complement the list of main rivalries with additional derbies mentioned in Andres et al. (2023) and on the German football rivalry Wikipedia page. Table ?? outlines the derby pairings and the respective data sources.

Mapping between datasets. All football matches pertaining to the placebo analysis contain a time variable and a geographic location - the coordinates of the venue where the match took place. We map the Cartesian coordinate of each venue to a German map, which contains polygons for each municipality. Merging to all other datasets leverages the unique municipality identifiers.

Table C.1. Description of Data Sources

Variable	Regional Level	Period	Source
Main variables			
Participants in PEGIDA protests	muni	2015-2020	Kanol & Knoesel (2021)
Hate crimes	muni	2015-2020	Amadeu Anotonio Foundation and PRO ASYL Foundation
Weather	muni	2015-2019	European Centre for Medium-Range Weather Forecasts (ECMWF)
Base controls			
GDP per capita	dist	2000-2019	Federal Statistics Office
Population density	muni	2009-2021*	own calculations
Unemployment rate	muni	2008-2021**	Federal Employment Agency
AfD votes (Bundestag)	muni	2013-2021***	Federal Returning Officer
Total crime cases (per 100k pop)	muni	2013-2021	Federal Criminal Police Office
Demographic controls			
Workers without qualification	dist	2008-2021****	Federal Statistics Office
Age distribution (0-25, 25-50, 50-75)	muni	2009-2021*	Federal Statistics Office
Gender distribution	muni	2009-2021*	Federal Statistics Office
Immigration-related controls			
Unemployment rate of non-Germans	muni	2008-2021**	Federal Employment Agency
Foreign workers with academic qualification	dist	2008-2021****	Federal Statistics Office
Asylum recipients (share)	dist	2011-2021*	Federal Statistics Office
Social media controls			
Twitter usage per capita	dist	baseline	Twitter API
#RefugeesWelcome tweets per capita	dist	baseline	Twitter API
AfD followers	muni group	baseline	Müller & Schwarz (2023)
Additional variables			
Tweets mentioning PEGIDA	dist	2015-2020	Twitter API

Note: This table provides information of the variables we use for the analysis. The first column describes the geographical level at which we observe each variables. Districts (402 of them) are equivalent to NUTS3, Municipalities are smaller. Column 2 describes the period for which we have information of each variable. * stands for up until 31.12 of the previous year; ** indicates that we have the information on the yearly/monthly average. *** during this period, every election **** as of 30.06 of the year. Column 3 provides the source from which we extract each of the variables.

Table C.2. **Diffusion: geographic and (right-wing) social media proximity to other large protest locations - using followers for PEGIDA**

	Hate crime in t	
	(1)	(2)
Log(participants)	0.0216*** (0.00514)	0.0215*** (0.00501)
PEGIDA social media prox. S_{it}^P	0.000266** (0.000111)	0.000174 (0.000108)
Overall social media prox. S_{it}^W		0.00140*** (0.000279)
Geographic prox. G_{it}		0.000102 (6.57e-05)
Observations	2,771,167	2,771,167
Municipalities	10,677	10,677
Adj. R-squared	0.092	0.093
Mean dep. var.	0.00282	0.00282
Week FE	Yes	Yes
Municipality \times MoY FE	Yes	Yes
Municipality linear time trend	Yes	Yes
All controls	Yes	Yes

Note: Differences in differences estimation with week and municipality-month of year fixed effects and municipality linear time trends. Full sample of municipalities. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01. Time horizon is January 2015 until December 2019. Controls include GDP per capita; population density; unemployment share; and lagged number of participants, protest and hate crimes, refugee share, vote share for the AfD in the latest European or national election and crime rate per 100K inhabitants respectively, all measured at the end of each calendar year. S_{it} is the social media network proximity-weighted sum of protest participants in other municipalities (standardized). Superscript P indicates PEGIDA network based on follows, superscript W indicators wider network, based on random sample of tweets. G_{it} geographic proximity weighted (linearly decreasing window function until 100km) sum of participants in other municipalities (standardized). Social media proximity is measured as the per capita number of retweets in municipality i over the previous six months of sampled tweets coming from municipality j . Outcome in columns 1 to 4 is any hate crime in the same week. Outcome in column 5 is any hate crime in the following week. Outcomes in columns 6 and 7 are whether a protest occurred on the same day and the log number of participants, respectively.

Table C.3. **PEGIDA Follows predict RTs, RTs predict follows**

	Followers of PEGIDA followers	PEGIDA retweets
	(1)	(2)
PEGIDA RTs	0.0552*** (0.0000825)	
Followers of PEGIDA followers		0.502*** (0.00750)
R ²	0.26	0.17
Observations	158,006	158,006
Source NUTS3 FE	Yes	Yes
Dest. NUTS3 FE	Yes	Yes

Note: Followers and RTs normalized by dividing by the population of both source and destination NUTS-3.