



ROCKWOOL Foundation Berlin

Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES

073/26

Right-Wing Protest and Hate Crimes

Sulin Sardoschau, Annalí Casanueva-Artís

Right-Wing Protest and Hate Crimes

Authors

Sulin Sardoschau, Annalí Casanueva-Artís

Reference

JEL Codes: D74, D72, K42, J15, L82

Keywords: Populism, hate crimes, protest, media

Recommended Citation: Sulin Sardoschau, Annalí Casanueva-Artís (2026): Right-Wing Protest and Hate Crimes. RFBerlin Discussion Paper No. 073/26

Access

Papers can be downloaded free of charge from the RFBerlin website: <https://www.rfberlin.com/discussion-papers>

Discussion Papers of RFBerlin are indexed on RePEc: <https://ideas.repec.org/s/crm/wpaper.html>

Disclaimer

Opinions and views expressed in this paper are those of the author(s) and not those of RFBerlin. Research disseminated in this discussion paper series may include views on policy, but RFBerlin takes no institutional policy positions. RFBerlin is an independent research institute.

RFBerlin Discussion Papers often represent preliminary or incomplete work and have not been peer-reviewed. Citation and use of research disseminated in this series should take into account the provisional nature of the work. Discussion papers are shared to encourage feedback and foster academic discussion.

All materials were provided by the authors, who are responsible for proper attribution and rights clearance. While every effort has been made to ensure proper attribution and accuracy, should any issues arise regarding authorship, citation, or rights, please contact RFBerlin to request a correction.

These materials may not be used for the development or training of artificial intelligence systems.

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com



Right-Wing Protest and Hate Crimes

Sulin Sardoschau^{Ⓔ*}

Annalí Casanueva-Artís^{Ⓔ†}

March 2026

Abstract

This paper examines the impact of right-wing populist mobilization on anti-minority violence across 10,000 German municipalities between 2015 and 2019. Exploiting variation in weather conditions on scheduled protest days, we show that right-wing protests held on pleasant days increase their salience and visibility by attracting larger crowds, generating more attention in traditional and social media and subsequently raising the probability of hate crimes by 8.6 percentage points. These offenses are carried out predominantly by known, recidivist, lone-actor extremists in the aftermath of the protest. Spillovers are substantial: downstream newspaper coverage of protests and social media networks transmit violence to municipalities that did not host any protest. Our findings highlight a critical externality of grass-roots populist movements: they not only drive immediate local violence but also propagate it across wider networks.

Keywords: Populism, hate crimes, protest, media

JEL classification: D74, D72, K42, J15, L82

*Humboldt Universität zu Berlin, RFBerlin and IZA. Email: sulin.sardoschau@hu-berlin.de

†Ifo Institute, Junior Research Group on Social Media and Center for Industrial Organization and New Technologies and CESifo. Email: casanuevaartis@ifo.de

Authors are listed in randomized order. We would like to thank Davide Cantoni, Ernesto Dal Bó, Ruben Durante, Oliver Falck, Vasiliki Fouka, Stephan Heblich, Horacio Larreguy, Ro'ee Levy, Alexey Makarin, Benjamin Marx, Federico Masera, Mauricio Romero, Carlo Schwarz, 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, 7th Monash-Warwick-Zurich Text-as-Data Workshop, POLECONUK Webinar, EEA-ESEM annual conference, Silvaplan Political Economics Workshop, Center for Economic Research Seminar at ITAM, CEMIR Junior Economist Workshop on Migration Research, The Political Economy of Democracy and Dictatorship (PEDD) Conference, CReAM/RWI Workshop on the Economics of Migration, VfS annual conference, participants of Junior Applied Microeconomists in Europe Seminar, the Rockwool Berlin Migration Forum, the Research Seminar Hohenheim/Stuttgart, Louis-André Gérard-Varet (LAGV), Public Economics Conference Marseille, 8th IBEO Workshop Alghero, Public Governance Working Group PSL, and the NYU Development Economics Seminar.

1 Introduction

Populist mobilization has become a fixture in Western democracies, with demonstrations against immigration, globalization, and perceived elite betrayal increasingly turning violent (ACLED, 2024; CSIS, 2022). An extensive literature has examined how populist actors influence economic policy, institutions, and social norms once in office (Guriev & Papaioannou, 2022; Funke et al., 2023). Yet, we know comparatively little about the impact of grassroots populist movements, and particularly their role in fueling politically motivated violence (Cantoni et al., 2024).

We study this issue in the context of Germany’s largest far-right movement since World War II, the Patriotic Europeans Against the Islamization of the Occident (PEGIDA). During Europe’s refugee crisis, PEGIDA and its local chapters organized regular anti-immigration demonstrations across Germany between 2015 and 2020. We exploit the fact that these protests were scheduled in advance and compare municipalities with a planned PEGIDA protest that experienced different local weather conditions on the protest day.¹ We estimate a difference-in-differences model with municipality and week fixed effects that interacts an indicator for a scheduled protest with an indicator for “pleasant” weather, controlling for each component separately. This interaction isolates weather-induced variation in *protest salience*: pleasant weather on a scheduled protest day increases turnout and media coverage on Twitter and in newspapers, and it increases the probability of a hate crime in the days following the protest by about 8.6 percentage points.²

Our empirical strategy relies on the assumption that it is as good as random whether a pre-scheduled PEGIDA protest takes place on a pleasant versus unpleasant day and influences subsequent hate crime only through protest salience. To probe the plausibility of this assumption, we conduct several exercises. First, we rule out strategic foresight by organizers, showing that pleasant weather does not predict the scheduling (or cancellation) of PEGIDA Monday protests. Second, we show that pleasant weather does not directly predict the incidence of hate crimes and verify that our results do not stem from serial correlation in weather conditions.³ Third, event studies reveal that pleasant-day protest locations are not on differential hate crime trajectories, suggesting no anticipation or intertemporal substitution effects.

To characterize the type of offenses driving the increase in hate crime incidence, we use a large language model (LLM) to classify hate crime descriptions, extracting details on perpetrator backgrounds, the number of perpetrators involved, crime locations, and whether incidents explicitly occurred in the context of a protest. The surge in hate crimes cannot be explained by protest-related incidents that include rallies, confrontations with counter-protesters or hate

¹Organizers drew on the legacy of the 1989 *Montagsproteste* in the former GDR and typically announced protest dates weeks in advance. Under German assembly law (*Versammlungsgesetz*), public demonstrations must be registered beforehand, so the timing of Monday protests is predetermined and plausibly orthogonal to short-run local conditions. Related work uses weather shocks to study participation in single, ad hoc protest events (Madestam et al., 2013; Beraja et al., 2023). In contrast, the repeated, pre-scheduled structure of PEGIDA demonstrations generates many local replications of the same design in a panel setting and focusing on the intensive margin.

²In our view, protest *salience* is plausibly the policy-relevant dimension since marginal participants and journalists may be more likely to respond to information about the potential negative externalities of attendance and coverage. However, we offer a complementary approach in Appendix A.1 that studies the extensive margin, leveraging quasi random timing in Islamic terrorist attacks outside of Germany as a protest trigger.

³If anything, it decreases the probability of hate crimes in the absence of protest. This is in line with the literature that links extreme (rather than moderate) weather conditions to criminal activity (Field, 1992; Heilmann et al., 2021). Our results also remain intact when controlling for weather conditions on each day of the week.

speech. Instead, the increase is concentrated in the days following the protest, and is driven by single-actor incidents and repeat offenders with a known extremist background, suggesting that protest salience activates individual-level anti-minority sentiment rather than facilitating the coordination of violence locally.

Motivated by these findings, we assess whether protest salience also increases hate crimes beyond the host municipality. If pleasant-day protests amplify the visibility of PEGIDA and its anti-minority narratives, their effects may propagate across municipal borders. We focus on two diffusion channels: geographic proximity and network exposure on social media. We measure geographic exposure to protest-hosting locations using a linear decay function up to 100km. We measure PEGIDA-specific social media connections by tracking users who retweet posts containing the term PEGIDA from accounts located elsewhere, lagged by six months. This measure captures the extent to which users in one municipality engage with PEGIDA-related content produced in another, reflecting ideological alignment, attention, or offline ties. To benchmark against general social media connectedness, we construct an analogous exposure measure using a random sample of tweets.

This approach allows us to examine the following scenario. Consider two municipalities, i and j , neither of which hosts a PEGIDA protest in week t . Both are equidistant to a third municipality l and have the same level of general social media exposure to l , but municipality i exhibits stronger PEGIDA-specific engagement with content produced in l . When l hosts a pleasant-day protest, do hate crimes increase more in i than in j ? We find that PEGIDA-specific social media ties predict higher hate crime incidence in municipalities that did not host a protest, and that this effect persists into the subsequent week. By contrast, the effect of broader social media connectedness and geographic proximity to pleasant-day protests is close to zero and insignificant. Together, these results indicate that pleasant-day protests diffuse primarily through PEGIDA-specific networks and operate by intensifying salience among already receptive audiences elsewhere.

Next, we examine the effects of media coverage of pro- and anti-PEGIDA protests directly. We proceed in two steps. First, we test whether pleasant-day PEGIDA protests increase the salience of local counter-mobilization. Greater PEGIDA salience could crowd out attention to opposition, but it could also raise the salience of the refugee issue more broadly and thereby elevate counter-narratives. Using newspaper articles from the GENIOS archive, we identify and geolocate mentions of counter-demonstrations. On average, each pleasant-day PEGIDA protest is associated with a 12 percentage point higher probability that a counter-protest in the same municipality and week is mentioned in any German newspaper. We also find a small increase in geolocated pro-refugee tweets.

Second, we test whether media exposure to PEGIDA protests and counter-protests affects hate crimes in municipalities that did not host protests. While our previous exercise on diffusion tests whether pleasant-day protests affect municipalities with pre-existing ties to protest locations, this exercise asks whether media exposure, independent of geographic distance or social media ties, shifts anti-minority violence elsewhere. To construct municipality-level newspaper exposure, we use the universe of GENIOS articles, count the number of articles published about PEGIDA in the previous week across all outlets in Germany, and weight this aggre-

gate time-varying newspaper exposure by the local circulation of each outlet from advertising statistics. We distinguish between articles reporting on PEGIDA-organized protests and articles on counter-protests. To construct an analogous measure of social media exposure, we measure the country-wide volume of tweets mentioning PEGIDA and the pro-refugee hashtag `#refugeeswelcome` over the previous week, weighting by the local Twitter users per capita at baseline.

Exposure to pro-PEGIDA coverage in print and online increases hate crimes: a one standard deviation increase in exposure to newspaper articles covering a PEGIDA protest is associated with a 0.04 percentage point higher probability of a local hate crime. In contrast, exposure to newspaper coverage of counter-protests has a mitigating effect: an equivalent increase is associated with a 0.02 percentage point decline. Pro-refugee tweets have no detectable effect, potentially because prospective offenders are not exposed to counter-narratives on social media due to echo-chambers and network effects. These results imply that downstream media coverage of right-wing protests propagates anti-minority violence to non-host municipalities, and that newspaper coverage of counter-protests partially offsets this diffusion.

Taken together, our findings reveal a previously under-appreciated externality of right-wing populist protests with implications for democracies worldwide grappling with rising extremism and anti-minority violence. The salience of right-wing protests triggers immediate local violence (potentially by shifting perceived norms and emboldening potential offenders), and these effects propagate through social media networks and downstream media coverage. Policy debates often focus on banning such assemblies altogether. Yet in settings where freedom of assembly protects right-wing mobilization, our results point to a more actionable lever: limiting the amplification of protest salience. In particular, media coverage and engagement by marginal participants can magnify the negative consequences, while counter-mobilization covered in mainstream outlets can partially offset them.

We situate our study within three strands of the literature. First, we add to prior work which shows that protests can shape mainstream political outcomes, often boosting electoral support (Madestam et al., 2013; González, 2020; Larreboure & González, 2021; Lagios et al., 2025) or swaying public opinion aligned with (Casanueva, 2021; Gethin & Pons, 2024; Cantoni et al., 2019) or in some cases opposed to the protesters' agenda (Caesmann et al., 2021) — all within the realm of the democratic process. A related strand of the literature examines how (social media) networks contribute to the spread of protests themselves (Qin et al., 2024; Lagios et al., 2025; Flückiger & Ludwig, 2025). To the best of our knowledge, we provide the first evidence that far-right protests increase hate crimes and that the social consequences of these protests (rather than mobilization itself) diffuse across space.

Second, we contribute to the literature on the societal consequences of populism. Existing work has documented how populist governments and parties erode macroeconomic stability, increase inflation and decrease real wages (Funke et al., 2023; Guriev & Papaioannou, 2022; Levitsky & Ziblatt, 2019; Edwards, 2019). Another strand of research highlights how populist victories reshape the media environment, activate anti-democratic ideologies and weaken democratic institutions (Funke et al., 2023; Guriev & Papaioannou, 2022; Bursztyjn et al., 2020). Similarly, signals from populist leaders can normalize discriminatory views (Bursztyjn 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; Bursztyrn & Yang, 2022). We contribute to this literature by showing how street-level populist mobilization, independent of formal political power, can foster anti-minority violence.

Third, 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; Bursztyrn et al., 2022), as well as to moments that shift societal norms (Hanes & Machin, 2014; Romarri, 2020; Carr et al., 2020; Bursztyrn et al., 2023). In some cases, economic and cultural factors interact. 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; Jiménez Durán et al., 2022). Our findings connect these two domains, documenting that social media networks act as powerful precursors to politically motivated violence.

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.

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.

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 late 2015. 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 referenced these demonstrations in an effort to gain legitimacy and pursue their stated goal of protecting 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 consolidated as a nationalist, xenophobic, and Islamophobic protest formation with recurrent overlap with neo-Nazi milieus (Vorländer et al., 2018; Virchow, 2016; Haller, 2019; Reuters, 2015; Connolly, 2014; Time Magazine, 2015). Its mobilization dovetailed with the AfD’s 2016 electoral breakthrough (Grabow, 2016; Weisskircher, 2023; Reuters, 2016b; Washington Post, 2016). In May 2021, Saxony’s domestic intelligence service classified PEGIDA as a proven extremist endeavor (Landesamt für Verfassungsschutz Sachsen, 2021; Associated Press, 2021). Bachmann’s record includes a conviction for incitement of hate at PEGIDA rallies (Reuters, 2016a; U.S. Department of State, 2020).

2.2 Main Data Sources

Protest. We take information on PEGIDA protests from the *Right-Wing Extremist Mobilization in Germany Data Set* (Kanol & Knoesel, 2021). The dataset is based on information from the German federal government’s responses to parliamentary questions tabled by the opposition left-wing party *Die Linke* and includes information on the location, date, number of participants, organizing actors, and motto of demonstrations classified as right-wing.

For our purposes, we restrict attention to Monday protests organized by PEGIDA or any of their local chapters and partnering organizations (including LEGIDA, BÄRGIDA, Die Rechte and others) because of their ritualized nature and because these were explicitly anti-Islam. Our final sample includes 475 protests between 2015 and 2020 with an average of 150 and a maximum of 2,300 participants (see Table A.2). However, it is worth noting that participation is measured with noise, based on rough police estimates and typically rounded up or down.⁴ Many of the protests were located in the former Eastern Germany but we show in Figure 1 that there are also offshoots in Western Germany.

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 January 2015 to January 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 antisemitic violence.⁵ 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 submits a parliamentary inquiry (*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

⁴There is substantial heaping at rounding intervals of 50 and 100 participants. Overall, there are only 12 observations where the number of participants is not a multiple of 5 or 10.

⁵This data has been used in other studies, including in Müller & Schwarz (2021) and Jiménez Durán et al. (2022)

case, the government reports its type (physical assault, arson and Neo-Nazi rally),⁶ date and location. It is worth noting that the vast majority of recorded hate crimes in our data set are related to (attempted or actual) physical assault or arson (typically of refugee camps), rather than hate speech. For a sub-sample of hate crimes, we have detailed descriptions of perpetrator, victim and context. We provide examples of these hate crimes in Appendix C.

Overall, there are more than 9,000 recorded hate crimes in the observation period. We show in Figure 1 that the cumulative number of hate-crimes per 100K inhabitants across municipalities between 2015 and 2020 is higher in the former East of Germany but prevalent throughout 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.

In Table A.3, we compare the characteristics PEGIDA-hosting municipalities to other municipalities. As expected, locations that host right-wing protest exhibit higher levels of violence against minorities: roughly 9% of municipality-week observations feature at least one hate crime, indicating a 35 times higher hate crime incidence. PEGIDA-hosting municipalities are also more densely populated, have a higher unemployment rate, a higher vote-share for the right-wing party, a lower overall crime rate and host a similar number of refugees per capita.⁷

Social Media. We use three measures to proxy social media use at the local level: *i*) overall Twitter use, *ii*) PEGIDA tweets, *iii*) 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.⁸ 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 proxy pro-refugee sentiment, using all tweets and retweets in German mentioning the hashtag #RefugeesWelcome between 2013 and 2018.⁹ We are able to geo-localize 150,000 of about 390,000 tweets.

Newspapers. The GENIOS newspaper database is a comprehensive digital repository offering full-text access to over 300 newspapers, including 180 German-language titles. This dataset encompasses the near-complete collection of print newspaper articles in Germany. To construct our dataset of newspaper coverage of pro-PEGIDA and anti-PEGIDA protest, we first filter for all articles mentioning PEGIDA, related organizations, keywords related to protests and migration, between 2015 and 2019. We then employ a large language model to analyze these

⁶We show examples of each type in Table C.2.

⁷We show in Appendix B that our results hold when restricting the sample to these municipalities.

⁸We 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.

⁹This data is taken from Jaschke et al. (2022) and thus only available for those years. Recent API restrictions imposed by Twitter/X do not allow us to extend this time frame.

selected articles, extracting information about protests and counter-protests mentioned therein. The construction process is detailed in Appendix section C.6.

Additional data sources. Appendix Table C.1 describes all regional controls, their geographic granularity, time coverage and frequency, as well as their sources. Regional controls come from four administrative sources. Labor market data are 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). 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 the Federal Statistics Office. We use NUTS3 boundaries of 2013 to harmonize administrative changes over time, which we describe in more detail in Appendix C.

3 Research Design and Main Results

3.1 Empirical Strategy

We study how the salience of right-wing protests, reflected in higher turnout and greater media coverage, affects anti-minority violence. A key concern is that protest salience may be correlated with unobserved, time-varying factors that also shape local hostility toward minorities. For example, communities with a stronger capacity or inclination for far-right mobilization may simultaneously attract larger crowds and experience higher levels of hate crime. In this case, we would confound the causal effect of protest salience with changes in underlying mobilization potential. To address this concern, we exploit exogenous variation in local weather conditions on predetermined protest days, which has been shown to impact political mobilization and media coverage (Gomez et al., 2007; Zhong & Zhou, 2012; Madestam et al., 2013; Jiang et al., 2022; Beraja et al., 2023).

These studies typically focus on a single, nationally coordinated protest day and use local weather variation to predict protest occurrence across space. In contrast, the German Assembly Act (*Versammlungsgesetz*) requires public assemblies to be registered in advance, making spontaneous or ad-hoc protests rare and limiting the scope for short-run shocks such as weather to affect protest occurrence. We therefore adopt a similar underlying logic in a setting where protests are repeated, pre-scheduled, and locally organized. Rather than comparing locations with pleasant versus unpleasant weather on a single day, we compare protests held under different weather conditions across municipalities and weeks. This yields a rich panel design and allows us to focus on the intensive margin of protest activity. Using data on more than 10,000 municipalities over 260 weeks between 2015 and 2019, we estimate a linear probability model with municipality and week fixed effects:

$$H_{it} = \alpha_i + \lambda_t + \gamma_1 (\text{Weather}_{it} \times \text{Protest}_{it}) + \gamma_2 \text{Protest}_{it} + \gamma_3 \text{Weather}_{it} + \tilde{\pi}'_{it} L_{it} + \nu_{it}, \quad (1)$$

$$\text{Weather}_{it} = \begin{cases} 1, & \text{if } \text{rain}_{it} \leq 10 \text{ mm/m}^2 \quad \& \quad 0^\circ\text{C} \leq \text{temp}_{it} \leq 20^\circ\text{C} \\ 0, & \text{otherwise.} \end{cases}$$

The outcome of interest, H_{it} , is a binary variable equal to one if at least one hate crime against refugees was recorded in municipality i during week t , and zero otherwise. Importantly, we exclude any hate crime recorded on the protest day itself, to distinguish between crimes committed after rather than during the protest itself.

The variable Protest_{it} indicates whether a Monday PEGIDA protest took place in municipality i in week t . As outlined in Section 2, these protests were the signature event of the movement, announced in advance and held regularly by local PEGIDA chapters and affiliated groups with the aim to invoke parallels to the peaceful civil rights protests of 1989 (*Montagsdemonstrationen*) that also took place regularly and on Mondays.¹⁰ We include week fixed effects λ_t and municipality fixed effects α_i , as well as a vector of controls $\tilde{\pi}'_{it}L_{it}$.¹¹ Standard errors are clustered at the municipality level to account for within municipality correlation over time. We show later that our results are not sensitive to alternative ways of clustering, including spatial clustering and two-way clustering at higher levels of aggregation.

We define “weather” as dummy variable for *pleasant weather conditions* that switches on at moderate temperatures (between 0°C and 20°C) and no heavy rain (precipitation lower than 10 mm per square meter as defined by the World Meteorological Organization) in the hours between 2 pm and 6 pm on the protest day, i.e. the time window shortly before and at the beginning of PEGIDA protests when participation decisions are plausibly made. This definition maximizes the predictive power for turnout as illustrated in Figure A.1, which shows the first-stage F-statistics for various cut-off combinations. As we will show below, our results are not sensitive to changes in these cut-offs. According to our preferred definition of the weather dummy, about one quarter of all protests (125 out of 475) fall into the unpleasant category and the rest into the pleasant category.¹² The time elapsed between two pleasant-day protests in the same municipality is about 6 weeks on average.

Our main specification accounts for time varying factors that are common to all municipalities, such as the overall popularity of the right-wing movement or salience of the refugee issue, as well as any unobserved time-invariant heterogeneity at the municipality level, including, for instance, the root determinants of anti-immigrant sentiment with time and municipality fixed effects. It also addresses concerns related to serial correlation in weather conditions and strategic timing by potential protesters. For instance, pleasant weather on a protest day in the preceding week could either encourage (riling up) or discourage (fatigue) attendance at future

¹⁰Outdoor assemblies in Germany are generally subject to a notification (Anzeige) requirement that is keyed to the time of public announcement rather than the event date. Thus, protests need to be announced at least 48 to 72 hours but since Monday protests belonged to the ritualized event, these protests were typically announced several weeks ahead. This is why we focus on Monday protest rather than any right-wing protest.

¹¹Recent 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 (e.g., De Chaisemartin & D’Haultfœuille (2023), 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 B.

¹²Figure A.2 shows temperature ranges and average precipitation. Figure A.2a shows that minimum and maximum temperatures range from -6°C to 31°C and documents the precipitation distribution with a large share of protests falling below 10 mm per square meter – the World Meteorological Organization’s (WMO) definition of heavy rain.

events. Alternatively, protesters may selectively attend based on weather forecasts. To control for such dynamics, our specification includes the control vector $\tilde{\pi}'_{it}L_{it}$ which accounts for lagged weather, protest, and pleasant-day protest, along with their interactions with contemporaneous protest and weather conditions.¹³

The coefficient of interest γ_1 captures the effect of a PEGIDA Monday protest held under pleasant weather relative to those held under unpleasant weather. In this model, γ_2 captures the association between protest weeks and hate crimes under unpleasant weather conditions (Weather = 0), while γ_3 captures the association between pleasant weather and hate crimes in non-protest weeks (Protest = 0). A causal interpretation of γ_1 relies on the assumption that, conditional on fixed effects and controls, it is as good as random whether the protest took place on a pleasant versus unpleasant day and that pleasant-day protests do not affect hate crimes through any other channels but protest salience.¹⁴

We dedicate Sections 3.2 and 3.3 to probe the plausibility of the identifying assumption in much detail – particularly the concern that pleasant weather impacts hate crimes directly – but let us highlight a few key points here. First, we show that there is no strategic scheduling of protests. In particular, we show that protest occurrence is orthogonal to local weather conditions, consistent with the idea that protests in Germany need to be scheduled and registered in advance and can thus not respond to short-run variations in weather conditions. Second, we check whether pleasant weather has a direct impact on hate crimes and show that the correlation is precisely estimated and negative (but close to zero), consistent with the literature that links extreme rather than moderate weather conditions to crime (Heilmann et al., 2021; Field, 1992). In addition, our results increase in magnitude and precision when controlling for weather on each day between Tuesday and Sunday of the protest week when hate crimes are recorded. Third, we show in an event study that municipalities with pleasant versus unpleasant Monday protests do not differ in their hate-crime trajectories prior to the protest, assuaging remaining concerns about strategic scheduling, anticipation effects or intertemporal substitution between pleasant-day protests and hate crimes.

Our main research design identifies the effect of protest *salience* rather than protest occurrence which is — in our view — the more policy-relevant dimension, given that freedom of assembly is constitutionally protected and cannot be curtailed lightly, whereas the intensity of protest activity, including turnout, media visibility, and broader communication, is shaped by individual and organizational choices. Nevertheless, the direct effect of protest occurrence remains of interest, particularly in debates where policymakers consider restricting or banning specific forms of protest altogether. We dedicate Appendix A.1 to study this question in more detail, leveraging plausibly exogenous variation in protest occurrence. Our results are consistent with and larger in magnitude than estimates of the average marginal effect of protest occurrence derived from our main design.

¹³Specifically, we add a vector of lagged controls (L_{it}) including previous week’s protest, weather, and their interaction, as well as interactions with current conditions. Formally, L_{it} comprises $\text{Protest}_{i,t-1}$, $\text{Weather}_{i,t-1}$, and $\text{Protest}_{i,t-1} \times \text{Weather}_{i,t-1}$ and $\tilde{\pi}_{it} = \pi + \phi \text{Protest}_{it} + \psi \text{Weather}_{it}$

¹⁴Following Borusyak & Hull (2023), we argue that one can obtain causal estimates even when exposure (who protests where/when) is non-random. Focusing on pre-determined protest days, controlling for protest and weather separately and including municipality and week fixed effects already address key concerns in this design. Moreover, we provide several additional plausibility checks proposed below.

3.2 Main Results

Protest Salience and Hate Crimes. Table 1 presents our main results. We report the coefficient for the interaction term of pleasant weather and protest, as well as coefficients for each term separately. At the bottom of the table, we also report the average marginal effect of protest alone. We begin by examining whether Monday PEGIDA protests held under pleasant weather conditions attract greater attention than protests held during adverse weather conditions in the form of turnout. Consistent with the literature on weather conditions and political participation (Gomez et al., 2007; Fujiwara et al., 2016), column 1 shows that pleasant-day protests increase turnout by 47 percent, equivalent to approximately 0.35 standard deviations.

We next assess whether pleasant-day protest translates into greater visibility in traditional and social media. Increased coverage may arise because larger crowds attract broader public attention, because participants themselves amplify the event on social media, or because journalists are more likely to attend and report on protests held under favorable weather conditions. To measure newspaper coverage, we draw on the GENIOS database, which contains the near-universe of articles published by almost 300 national and local German outlets. We first restrict attention to articles mentioning PEGIDA or affiliated organizations, yielding 142,707 articles. We then apply a large language model to identify coverage of specific protest events and their locations (we describe this procedure in more detail in Appendix C). This allows us to match 2,188 articles to 82 protests in our dataset. Column 2 shows that PEGIDA protests held under pleasant weather generate a 34 percent higher newspaper coverage than protest held under unpleasant weather conditions.

We conduct a similar analysis for social media coverage using the universe of Tweets mentioning PEGIDA and geo-locating authors based on self-reported profile locations (again, we describe the procedure in more detail in Appendix C). The final sample comprises approximately 660,000 geo-located Tweets over three years. Analogous to the newspaper results, column 3 shows that pleasant-day protests are associated with a 23 percent increase in the number of geo-located Tweets mentioning PEGIDA.¹⁵ Together, these findings suggest that favorable weather substantially amplifies the salience of PEGIDA protests through increased turnout and media coverage. Such increases in salience may in turn contribute to downstream violence locally but also in other locations, consistent with Sabet et al. (2025), who document that increases in support for the right-wing party *AfD* following right-wing terrorist attacks in Germany are driven by media salience. We return to this idea in Section 4.

Finally, we turn to the effect of protest salience on hate crimes in column 4. Protests held under pleasant weather conditions increase the probability of recording a hate crime in the six days following the protest by 8.6 percentage points, indicating that higher protest salience substantially amplifies anti-minority violence. Table A.4 reports the same specification for the sub-sample of municipalities that hosted at least one PEGIDA Monday protest. These locations have a substantially higher baseline risk of anti-minority violence (see Table A.3). Restricting the sample to ever-treated municipalities leaves our coefficient of interest virtually unchanged in magnitude and precision, while the coefficient on protest occurrence itself becomes small and

¹⁵Table B.8 shows robustness of the results to alternative outcome definitions, including an inverse hyperbolic sine transformation, per municipality residents and standardization.

imprecise. This pattern is consistent with protests taking place in locations where anti-minority sentiment is already present, so the mere occurrence of a protest does little to shift behavior. What matters is variation in salience: when favorable weather increases the visibility and scale of the protest, anti-minority actors appear more willing to commit hate crimes. We present additional evidence consistent with this interpretation below.

Throughout, we focus on the reduced-form effect of pleasant-day protests on hate crimes because protest salience is multidimensional and the weather shock jointly shifts turnout, traditional media coverage, social-media activity, and potentially other channels of visibility. For interpretation, however, Table A.5 reports 2SLS estimates using participation alone and a composite salience index constructed as the first principal component of participation, newspaper coverage, and PEGIDA-related tweets. In both cases the first stage is strong, with Kleibergen–Paap F-statistics of 23.11 and 20.93. The estimates imply that a ten percent increase in turnout raises the probability of recording a hate crime in the days following the protest by 2.1 percentage points, while a one standard deviation increase in the composite salience index raises it by 1.1 percentage points.¹⁶

Plausibility of the Identifying Assumption. Table 2 reports a set of placebo and validation exercises that probe the plausibility of the identifying assumption. We examine whether pleasant weather directly affects protest occurrence or hate crimes, and whether persistent weather patterns or other right-wing protest activity confound our baseline estimates.

Public demonstrations in Germany must be registered in advance, implying that the timing of PEGIDA’s Monday protests is predetermined and orthogonal to realized local weather. A remaining concern is that adverse weather could lead to protest cancellations. Column 1 shows that this is not the case: pleasant weather does not predict the occurrence of a Monday protest. This confirms that weather shifts turnout and salience but not the extensive margin of protest activity. Column 2 tests for a direct effect of pleasant weather on hate crimes. Consistent with evidence that crime responds to extreme weather rather than moderate conditions (Field, 1992; Heilmann et al., 2021), the estimated coefficient is small, precisely estimated, and close to zero, ruling out a direct weather-crime channel.

Columns 3 and 4 re-estimate the main specification while conditioning on persistent weather conditions and on other right-wing protests occurring in the same week. Controlling for weather persistence increases both the precision and magnitude of the estimated effect in column 3. Column 4 shows that the estimate is unchanged when controlling for other right-wing protests at the municipality-week level, indicating that the baseline effect does not reflect heightened protest activity in the days surrounding the Monday protest.

Columns 5 and 6 directly test whether PEGIDA Monday protests held during pleasant weather predict the number or turnout of other right-wing protests. This addresses the concern

¹⁶The 2SLS and reduced-form coefficients are roughly an order of magnitude larger than the OLS estimate and are estimated more precisely. Several mechanisms can account for this difference. First, reported crowd sizes are measured with noise; classical measurement error attenuates the OLS toward zero, whereas the instrument relies on precisely measured protest timing and weather conditions. Second, treatment effects may be heterogeneous. The instrument identifies a local average treatment effect for weather-induced variation in salience, including turnout. OLS averages across all demonstrations, including predictable local support and turnout for PEGIDA that potential perpetrators may already anticipate and therefore incorporate into their behavior. As we show below, the estimated effect reflects salience shocks that embolden already radical or recidivist actors rather than marginal attendees whose participation depends on weather, which is consistent with a downward bias in OLS.

that pleasant-day protests may coincide with unobserved, time-varying surges in local mobilization potential or anti-minority sentiment. We find no such correlation. While the protest indicator itself is positively and precisely related to other protest activity, consistent with shared scheduling or underlying local conditions, pleasant-day protests are not.¹⁷ The small and imprecise coefficient in column 6 further indicates no meaningful complementarity or substitution in attendance between pleasant-day PEGIDA protests and other right-wing protests within the same week.

Hate-Crime Characteristics. Our main specification focuses on hate crimes committed in the days following the protest, thus excluding hate crimes committed in anticipation of or in the immediate aftermath. In this section, we assess the extent to which protest-related crimes contribute to the overall surge in hate crimes and examine whether pleasant-day protests facilitate coordination among prospective offenders and encourage newcomers to commit hate crimes.

To assess this possibility, we use a large language model (LLM) to classify hate crime descriptions, extracting details on perpetrator backgrounds, the number of perpetrators involved, crime locations, and whether incidents explicitly occurred in the context of a protest or rally. We manually verify the accuracy of this classification for a 10 percent random sample of hate crime descriptions.¹⁸ Figure 3 summarizes results from separate regressions, plotting the coefficient, 90 percent, and 95 percent confidence intervals for pleasant-day protests. The top row reproduces the baseline estimate as a reference.

The second row presents the coefficient after including hate crimes committed on the same day. The estimate remains precise and slightly higher in magnitude. This is in line with the literature that documents a surge in violence through pure crowd size and heightened emotions (Andres et al., 2023; Branscombe & Wann, 1992). Next, we focus explicitly on protest-related offenses, defined as those that reference protests, rallies, or public incitement to violence against minorities. These types of crimes increase in the days following the protest but do not account for the overall increase in hate crimes.

We then consider perpetrator characteristics, distinguishing between hate crimes committed by individuals with a documented record of right-wing extremism or prior criminal history (“recidivist perpetrators”) and those committed by individuals previously unknown to police. The increase in hate crimes is concentrated among repeat offenders. Finally, we examine whether offenses are committed by groups or individuals acting alone. While the coefficient for group-based crimes is imprecisely estimated, we detect a larger and significant increase in hate crimes perpetrated by lone actors.

Taken together, these findings suggest that coordination at protests contributes to the observed rise in anti-minority violence but cannot fully explain it. Pleasant-day protests may facilitate the recruitment of new or “marginal” offenders but the spike in hate crimes is most pronounced among known extremists and single perpetrators, pointing to an indirect channel whereby large protest turnout emboldens anti-minority violence among those most prone to it.

¹⁷The positive coefficient on the protest indicator in column 5 indicates that PEGIDA protests tend to coincide temporally with other right-wing demonstrations. Importantly, this pattern does not vary with weather conditions.

¹⁸Details on the classification methodology are provided in Appendix C.

3.3 Additional Validation and Robustness Checks

Anticipation and intertemporal substitution. As mentioned above, our preferred specification controls for lagged protest, weather conditions and pleasant-day protest as well as their respective interactions with contemporaneous weather conditions and protest occurrence to account for serial correlation in weather and protest intensity. In an event-study, we further examine the possibility of strategic timing and anticipation effects of scheduled protest and hate crimes, which would confound our estimates (Borusyak & Hull, 2023). Figure B.1 estimates the dynamic treatment effect of pleasant-day protest, including municipality and week fixed effects as well as protest and weather separately (but excluding the vector of lagged controls).¹⁹ Reassuringly, we detect neither a differential increase in violence nor an Ashenfelter dip in the six weeks leading up to a pleasant-day protest, assuaging concerns about strategic timing of pleasant-day protest, anticipation effects or inter-temporal substitution of hate crimes. In addition, we show in columns 1 and 2 of Table B.1 that our results hold when we include the cumulative number of past PEGIDA protests and the cumulative number of past participants, which addresses concerns about persistent or compounding effects of previous protest and turnout.

Seasonality. While common seasonal patterns across municipalities are already captured by week fixed effects, it is possible that seasonal patterns that are municipality-specific (such as higher summertime unemployment) correlate with both pleasant weather and hate crimes. Figure 4 illustrates the monthly distribution of protests and pleasant-day protests from January 2015 to December 2019. Pleasant-day protests are evenly distributed across the calendar, not clustering in any particular season. In addition, we show in column 3 of Table B.1 that our results hold when we include municipality-by-month-of-year fixed effects. This captures all municipality-specific seasonal variation, including aforementioned variables like unemployment but also typical weather conditions in the locality in a specific month. Thus, in this particular specification, the pleasant-weather indicator can be interpreted as deviations from average local weather conditions and thus re-centers the reduced form estimate (Borusyak & Hull, 2023). The coefficient of interest remains stable and precise, confirming that municipality-specific seasonal factors do not drive our results.

Unobserved municipal characteristics. Finally, we provide evidence supporting the plausibility of our identifying assumption that pleasant-day protests are uncorrelated with unobserved municipality characteristics. Figure B.2 plots standardized coefficients from regressions predicting the number of pleasant-day protests in a year with extensive municipal-level characteristics: demographic structure, socioeconomic status, political preferences, cultural composition, migration statistics, and crime data.²⁰ We observe no systematic relationship between

¹⁹We estimate $H_{it} = \alpha_i + \lambda_t + \sum_{\substack{k=T_0 \\ k \neq -1}}^{T_1} \beta_k (\text{weather}_{ik} \times \text{protest}_{ik}) + \gamma_k (\text{weather}_{ik}) + \phi_k (\text{protest}_{ik}) + \epsilon_{it}$ for the ever-treated sample but results hold for the full sample of municipalities.

²⁰Specifically, we estimate regressions of the form described by $\sum \text{Pleasant} \times \text{Protest}_{\text{muni.year}} = \beta_1 X_{\text{muni.year}} + \beta_2 \sum \text{Pleasant}_{\text{muni.year}} + \beta_3 \sum \text{Protest}_{\text{muni.year}} + \gamma_{\text{state}} + \theta_{\text{year}} + \epsilon_{\text{muni.year}}$, i.e. we regress the cumulative number of pleasant-day protest in a given municipality and year on an array of municipal characteristics, controlling for the total number of protest and pleasant days, and including state and year fixed effects and clustering standard errors at the state level.

these characteristics and the number of pleasant-day protests. Furthermore, our results remain unaffected when including a large set of municipality-level controls (Panel A of Table B.2) and when interacting them with protest and weather separately (Panel B of Table B.2), thus accounting for heterogeneous effects of protest occurrence and weather conditions in municipalities with different baseline characteristics.²¹

Bias in reporting of hate-crimes One potential concern with our analysis is differential reporting of hate crimes if, after pleasant-day 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. To empirically test for reporting bias, we conduct several exercises in Table B.3. Since 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 regional reporting changes.

Next, we examine whether police effort increased at the local level. 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 and thus less likely to suffer from reporting bias at the margin. Nonetheless, we examine the share of plausibly detectable hate crimes in column 5. Arson is typically more severe, almost exclusively targets refugee accommodations 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 a shift towards more sensitive crime reporting. Reassuringly, we find no evidence for a shift away from more easily observable hate crimes. Let us also preview here that our evidence on spatial diffusion patterns of hate crimes in the next section also speaks against reporting bias.

Additional robustness checks To ensure the reliability of our findings, we conduct a battery of robustness checks addressing several potential concerns, which we summarize briefly here and in more detail in Appendix B. First, we confirm that our results are not sensitive to the specific weather thresholds used to define pleasant-day protest. Figure B.3 shows that estimates remain stable across a range of plausible cut-offs for precipitation and temperature.

Second, we examine potential spatial spillovers. In Table B.4, we replicate our analysis at higher levels of geographic aggregation, including 401 NUTS-3 (column 2) and 38 NUTS-2 regions (column 3) in Germany. Results become larger and remain precisely estimated for these larger units. Columns 4 to 7 replicate the baseline specification but employ Conley standard errors with distance thresholds of 50 km, 100 km, 150 km, and 200 km to account for spatial correlation, confirming the robustness of our estimates. In addition, we provide

²¹In Panel A, we successively introduce GDP per capita, population density and unemployment rate (column 2), share of refugees in the municipality (column 3), vote share for the right-wing party AfD in the latest national or European election (column 4) and overall crime rate (column 5). Their respective interactions with pleasant weather and protest are included in Panel B of the same table.

alternative clustering approaches in Table B.5 which report two-way clustered standard errors by municipality and week, month and quarter, respectively, as well as two-way clustering by district and month, as well as Driscoll-Kraay standard errors.

Third, we introduce a more stringent set of fixed effects to account for potential unobserved time-varying municipality characteristics. Table B.6 shows that our results are robust to introducing municipality by month of year fixed effects (column 2), municipality specific linear time trends (column 3), both (column 4), as well as state-week fixed effects (column 5). As mentioned above, our results are robust to the inclusion of a large set of municipal controls as well as their interaction with weather and protest (see Table B.2).

Fourth, we examine the sensitivity of our results with respect to outliers. Figure B.4 estimates reduced form regressions dropping single NUTS-3 regions and single weeks at a time. Throughout, estimates are statistically indistinguishable from one another and remain precisely estimated. Finally, we test other estimation methods. Recent work in the two-way fixed effects literature emphasizes that multiple or staggered treatments can lead to biases (Roth et al., 2022). To address this concern, in Table B.7 we employ the estimator from De Chaisemartin & D’Haultfoeuille (2022) to estimate the treatment effect under the assumption that the effect stabilizes at within two weeks (column 2) or within six weeks (column 3) after the pleasant-day protest. We also show non-linear models in column 4 and 5, where we estimate a logit regression of the likelihood of observing a hate crime and a Poisson regression to estimate the count of hate crimes in the same municipality and week. Throughout, our results remain robust to changes in estimation methods. We also show that our results are largely robust to alternative definitions of the outcome variables, including the standardization, inverse hyperbolic sine transformation and per capita definitions in Table B.8.

4 Mechanisms

PEGIDA protests held under pleasant weather conditions attract larger crowds, receive more coverage in traditional and social media, and are followed by a surge in hate crimes. We interpret our reduced-form estimates as capturing the effect of protest salience. Salience need not operate locally: increased visibility may affect behavior beyond protest-hosting municipalities, depending on how attention propagates across space and networks. We study the diffusion of anti-minority violence to non-hosting locations through geographic proximity and social media exposure. We also examine whether pleasant-day PEGIDA protests raise the salience of counter-mobilization and whether media coverage of PEGIDA and counter-protests affects hate crimes outside protest locations.

4.1 Spatial Diffusion: Geographic Proximity and Social Media Networks

To assess the spatial diffusion of violence, we construct three measures of exposure to pleasant-day protests, focusing on municipalities that did not host a protest in a given week. First, we define geographic diffusion G_{it} as a distance-weighted measure of pleasant-day protests occurring in all other municipalities. Specifically, the exposure of municipality i to a pleasant-day protest in municipality j is calculated using the distance between municipal centers, applying a linear

decay function with a 100 kilometers cutoff. We then aggregate this bilateral measure across all municipalities and standardize it, which yields the geographic exposure G_{it} . More formally, if D_{ij} is the distance between i and j in kilometers, we define:

$$G_{it} = \sum_{j \neq i, D_{ij} \leq 100} \left(1 - \frac{D_{ij}}{100}\right) \times (\text{Protest}_{jt} \times \text{Weather}_{jt})$$

Second, we create two measures of digital diffusion based on Twitter networks rather than geographic distance, differentiating between PEGIDA-specific networks and general social media connections. Formally, we define digital diffusion as

$$S_{it}^{\Gamma} = \sum_{j \neq i} N_{j \rightarrow i(t)}^{\Gamma} \times (\text{Protest}_{jt} \times \text{Weather}_{jt}),$$

where Γ denotes either the PEGIDA network P or a broader Twitter network W . $N_{j \rightarrow i(t)}^{\Gamma}$ captures the extent to which users in location i engage with content produced in location j . For $N_{j \rightarrow i(t)}^P$, we leverage information on the universe of tweets and retweets that contain the keyword PEGIDA, identifying the original tweet’s user location and the retweeting user’s location. In particular, $N_{j \rightarrow i(t)}^P$ is constructed as the total number of PEGIDA tweets that were produced in j over the preceding six months and subsequently retweeted by users located in i divided by the population of i to get at the relative network size.

Importantly, this measure is directional: municipality j ’s influence on i may differ from its reciprocal, as retweet patterns are not necessarily symmetric. We interpret S_{it}^P primarily as a measure of local engagement with PEGIDA content produced elsewhere. Such engagement could reflect ideological alignment, real-world social ties, curiosity, or critical monitoring by opponents, but in any case implies heightened exposure to protest-related narratives.

In addition, we construct a similar measure based on a random sample of 600,000 tweets, capturing general Twitter connectivity between municipalities rather than PEGIDA networks specifically.²² This general measure quantifies overall Twitter interactions, reflecting how much users in municipality i engage with any content produced in municipality j . We again aggregate retweets over a six-month window and normalize by the receiving municipality’s population.

Analogous to our baseline specification, we estimate a linear probability model including municipality and week fixed effects. All diffusion measures are standardized for ease of comparison. We restrict the sample to municipalities that did not host a protest in a given week but still include weather and the interaction of weather with lagged variables (weather, protest, and pleasant protest) as controls. Standard errors are clustered at the municipality level. We estimate the following regression:

$$H_{it} = \beta_1 S_{it}^P + \beta_2 S_{it}^W + \beta_3 G_{it} + \beta_4 \text{Weather}_{it} + \tilde{\pi}'_{it} L_{it} + \delta_t + \mu_i + \epsilon_{it} \quad (2)$$

This allows us to document the diffusion patterns of anti-minority violence, examining whether locations that are geographically close (G_{it}), and digitally connected (S_{it}^P and S_{it}^W) to muni-

²²Appendix C provides further details on data collection and geo-location procedures. All Twitter data is matched at the NUTS-3 level, a higher level of geographic aggregation.

palities that host a PEGIDA protest on a pleasant day experience an increase in hate crime in the same week.

Table 3 presents the results. Columns 1 to 4 estimate the spatial spillovers of pleasant-day protests on hate crimes in the same week, while column 5 examines persistence by shifting the outcome to the following week. In column 1, we find a small and marginally significant effect of geographic proximity G_{it} . Column 2 considers exposure through general social media networks S_{it}^W , revealing a modest positive effect: a one standard deviation increase in exposure from general Twitter networks in a given municipality raises the probability of hate crimes by 0.042 percentage points. The coefficient for PEGIDA-specific networks S_{it}^P in column 3 is large and precisely estimated. A one standard deviation increase in PEGIDA network weighted exposure to pleasant-day protests significantly increases hate crime probability by 0.087 percentage points.

Column 4 includes all measures jointly. The coefficient on geographic proximity becomes insignificant, indicating that physical closeness alone does not generate spillovers in violence.²³ Similarly, the estimate for general social media proximity becomes negative and insignificant, suggesting that general social media exposure does not independently drive the diffusion of violence. In contrast, the effect of PEGIDA-specific networks remains robust and slightly increases to 0.09 percentage points. Column 5 further demonstrates persistence, showing that PEGIDA network exposure predicts hate crimes in the following week with a coefficient of 0.11 percentage points.

These results are consistent with a mechanism whereby increased salience emboldens individuals in locations with a pre-existing interest in the movement, leading to a surge in hate crimes even in municipalities without protests themselves.²⁴ These results complement previous work highlighting how online networks facilitate protest diffusion (Qin et al., 2024; Flückiger & Ludwig, 2025), while extending this insight to show that social media networks can also propagate the broader social consequences of a movement, beyond mobilization alone.

4.2 Media Coverage

Coverage of Counter-Protest. Pleasant-day PEGIDA protests may not only increase the salience of right-wing mobilization but also shape the visibility of opposing narratives. A priori, the direction of this effect is ambiguous. Heightened attention to PEGIDA may crowd out counter-narratives if media and online discourse focus disproportionately on the protest itself. Alternatively, greater overall attention to the issue may amplify coverage of counter-mobilization, raising the salience of both the movement and its opponents. We therefore study whether pleasant-day PEGIDA protests increase the coverage of counter-protests in newspapers and on social media.

To do so, we collect newspaper articles from GENIOS that reference anti-PEGIDA protests, extracting information on their location, date, and match them to municipality-weeks in our

²³The small and imprecisely estimated coefficient also suggests that spatially correlated weather conditions are unlikely to significantly drive anti-minority violence, and that the surge in hate crimes in protest-hosting locations reflects a level increase rather than a substitution away from nearby areas.

²⁴We show in Table A.6 that our results are very similar when including the protest hosting locations and estimate our baseline specification 1 with the full set of controls. There, PEGIDA networks to pleasant-day protest hosting locations influence local hate crimes above and beyond local pleasant-day protest.

data. Appendix C provides further details.²⁵ Columns 1 and 2 of Table 4 report estimates from our preferred specification but use as an outcome the likelihood that a counter-protest is covered in newspapers, measured in the same week as the PEGIDA protest and in the subsequent week. We find that pleasant-day PEGIDA protests significantly increase contemporaneous media coverage of counter-protests. Column 1 implies that a pleasant-day PEGIDA protest is associated with a 11.7 percentage point increase in the likelihood that the counter-protest is also covered. Column 2 suggests that this effect is short-lived and does not persist into the following week.

Next, we examine coverage online. To capture pro-refugee mobilization on social media, we collect geo-localized tweets containing the hashtag `#refugeeswelcome`. Manual validation of a random sample of 1,000 tweets confirms that more than 98 percent express support for asylum seekers in Germany. Regressing the volume of these tweets on the occurrence of pleasant-day PEGIDA protests, we find a significant increase of 32 percent in the number of pro-refugee tweets during the protest week, with no corresponding increase in the subsequent week. Taken together, these results show that salient PEGIDA protests amplify not only right-wing messages but also the visibility of counter-mobilization in both traditional and social media. These responses are immediate and sizable but fade quickly. In the next section, we assess whether exposure to these counter-narratives mitigates the increase in anti-minority violence following PEGIDA protests.

Exposure to Pro- and Anti-PEGIDA Coverage. Pleasant-day PEGIDA protests substantially raise salience, not only by increasing turnout but also by amplifying media coverage of both PEGIDA and counter-mobilization in newspapers and on social media. In this section, we examine whether local exposure to pro- and anti-PEGIDA coverage shapes anti-minority violence in non-hosting municipalities. Relative to the spatial diffusion exercise, which traces spillovers through geographic proximity and direct digital connections, we focus on a broader media-salience channel and ask whether exposure to PEGIDA and counter-PEGIDA narratives affects hate crime even in the absence of local protest activity or pre-existing ties to protest-hosting locations.

First, we focus on newspaper coverage. As before, we exploit data on the universe of German newspaper articles from GENIOS and combine this with municipality-level readership from advertising statistics (see Appendix C for more detail). We identify articles that mention PEGIDA protest as well as counter-demonstrations using a LLM. For each newspaper, we multiply the newspaper’s circulation in each municipality by the number of published articles referencing pro-PEGIDA (or anti-PEGIDA) demonstrations in the previous week, then sum over all newspaper outlets and divide by the municipality’s population to capture per capita exposure.

Similarly, we develop a time-varying measure of local social media exposure to PEGIDA content and pro-refugee content, using the aggregate number of tweets mentioning PEGIDA and `#refugeeswelcome` respectively. We proxy per capita exposure to Twitter content by the number of geo-localized Twitter users from a random sample of Tweets divided by the municipality’s population.²⁶ In essence, our indices capture the share of the local population

²⁵Note that there is no comprehensive data set on counter-protests collected by police authorities or the government.

²⁶Newspaper readership data is available for November 2019 and Twitter penetration is measured at the NUTS-3 level

exposed to aggregate PEGIDA content produced in newspapers and on Twitter over the previous week. Our exposure measures write as follows:

$$\text{Newspaper } \sigma_{i,t} = \frac{\sum_p \text{Articles PEGIDA}_{p,t}^\sigma \times \text{Readership}_{p,i}}{\text{Population}_i}$$

$$\text{Twitter } \sigma_{i,t} = \frac{\text{Tweets PEGIDA}_t^\sigma \times \text{Twitter Users}_i}{\text{Population}_i},$$

where $\sigma \in (\text{pro}, \text{anti})$ stands for pro and anti-PEGIDA coverage and p is an indicator for one of 286 newspaper outlets in Germany. We standardize both measures of exposure for ease of comparison and lag them by one week. Again, we estimate a linear probability with municipality and week fixed effects for the sample of non-hosting municipalities and cluster standard errors at the municipal level. In particular, we estimate:

$$H_{it} = \sum_{\sigma \in (\text{pro}; \text{anti})} \beta_\sigma \text{Newspaper } \sigma_{i,t-1} + \sum_{\sigma \in (\text{pro}; \text{anti})} \gamma_\sigma \text{Twitter } \sigma_{i,t-1} + \delta_t + \mu_i + \epsilon_{it} \quad (3)$$

This allows us to examine whether aggregate surges in newspaper coverage and Twitter discussions in the previous week impact anti-minority violence in places that are exposed to these news sources but did not host protests themselves. More explicitly, we assume that week-to-week variation in local newspaper and Twitter exposure to PEGIDA content is driven by factors orthogonal to unobserved drivers of hate crimes, after controlling for municipality and week fixed effects and restricting the sample to non-hosting locations.²⁷

Table 5 reports our results. We focus on social media exposure first. Column 1 reveals a strong and statistically significant positive relationship between exposure to PEGIDA content on Twitter and hate crimes in locations that did not host a protest in a given week. A one standard deviation increase in Twitter exposure to PEGIDA-related content in a municipality is associated with a 0.05 pp increase in the probability of observing at least one hate crime in the same week. If there were a global one standard deviation in pro-PEGIDA Twitter content across all municipalities, this would translate to 5.5 additional hate crimes. In contrast, the coefficient for pro-refugee content is not significant, suggesting that pro-refugee content does not affect hate crimes. A possible explanation is that this content does not reach the potential perpetrators because they do not follow accounts that would spread these messages.

Exposure to traditional media follows a similar pattern. Column 2 shows that municipalities more exposed to newspaper coverage of pro-PEGIDA protests experience a surge in hate crimes, whereas coverage of anti-PEGIDA protest is associated with a reduction in hate crimes. In particular, a one standard deviation increase in exposure to PEGIDA protest coverage in

using a random sample of tweets from 2014.

²⁷This logic is similar to a leave-out shift share instrument, where time variation comes from surges in coverage are driven by protest held elsewhere, which acts as an aggregate exogenous shifter in attention to PEGIDA and where cross-sectional variation comes from local readership or Twitter penetration, such that aggregate shifts disproportionately impact these locations.

newspapers increases the probability of a hate crime by 0.042 pp (or 4.5 hate crimes across all municipalities), while a one standard deviation increase in exposure to counter-protests decreases the hate crime probability by around 0.016 pp (1.7 hate crimes across municipalities). Column 3 includes all four exposure measures jointly. The results remain stable and of comparable magnitude. Column 4 confirms that the effects persist into the subsequent week, with similar magnitudes.

Overall, these findings point to a mechanism whereby exposure to PEGIDA-related content increases anti-minority violence in places that are not hosting any protest and that are not necessarily connected through physical and digital networks, potentially through an emboldening effect on prospective offenders elsewhere.²⁸ However, counter-narratives, at least when spread through traditional media, can mitigate this emboldening effect and reduce hate crimes.

5 Conclusion

This paper documents that right-wing mobilization bears substantial societal costs. We examine the relationship between grassroots populist mobilization and anti-minority violence in the context of Germany’s largest far-right movement since World War II, PEGIDA. Using plausibly exogenous variation in protest salience driven by weather conditions, we document that pleasant-day protests significantly raise the probability of hate crimes, particularly among lone-actor extremists and known perpetrators. We demonstrate that these effects extend beyond direct participants or immediate observers, propagating through broader media attention and social media networks. While PEGIDA-related media coverage plays a critical role in spreading protest-related narratives and subsequently inciting hate crimes, exposure to counter-protests partly mitigates these violent spillovers.

Our findings highlight an overlooked dimension of populist mobilization: even absent formal political power, grass-roots right-wing movements fuel politically motivated violence. Our results also add to research on the determinants of hate crimes, connecting economic and social factors with online extremism. We demonstrate how offline protests and online networks jointly amplify extremist messages, underscoring the importance of media environments in shaping these outcomes. Future research could further unpack the mechanisms behind media-driven radicalization and investigate long-term consequences of repeated populist mobilization.

²⁸In Table A.7, we repeat this exercises for the full sample of municipalities, estimating our preferred baseline specification (including interacted controls for protest and weather) and adding the media exposure measures. Our results remain similar in magnitude and show that media coverage of protest impacts hate crimes above and beyond local pleasant-day protest.

References

- ACLEDD (2024). Acled 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.
- Associated Press (2021). Saxony's domestic intelligence classifies pegida as extremist. AP News summary in English.
- 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).
- Berntzen, L. E. & Weisskircher, M. (2016). Anti-islamic pegida beyond germany: Explaining differences in mobilisation. *Journal of Intercultural Studies*, 37(6), 556–573.
- Borusyak, K. & Hull, P. (2023). Nonrandom exposure to exogenous shocks. *Econometrica*, 91(6), 2155–2185.
- Branscombe, N. R. & Wann, D. L. (1992). Role of identification with a group, arousal, categorization processes, and self-esteem in sports spectator aggression. *Human relations*, 45(10), 1013–1033.
- 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).
- Bursztyn, L., Egorov, G., & Fiorin, S. (2020). From extreme to mainstream: The erosion of social norms. *American Economic Review*, 110(11), 3522–3548.
- Bursztyn, L., Egorov, G., Haaland, I., Rao, A., & Roth, C. (2022). Scapegoating during crises. In *AEA Papers and Proceedings*, volume 112 (pp. 151–155): American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Bursztyn, L., Egorov, G., Haaland, I., Rao, A., & Roth, C. (2023). Justifying dissent. *The Quarterly Journal of Economics*, 138(3), 1403–1451.
- Bursztyn, L. & Yang, D. (2022). Misperceptions about others. *Annual Review of Economics*, 14, 425–452.
- Caesmann, M. J., Caprettini, B., Voth, H.-J., & Yanagizawa-Drott, D. (2021). *Going viral: propaganda, persuasion and polarization in 1932 Hamburg*. Technical Report 16356.
- Cantoni, D., Kao, A., Yang, D. Y., & Yuchtman, N. (2024). Protests. *Annual Review of Economics*, 16.
- Cantoni, D., Yang, D. Y., Yuchtman, N., & Zhang, Y. J. (2019). Protests as strategic games: experimental evidence from hong kong's antiauthoritarian movement. *The Quarterly Journal of Economics*, 134(2), 1021–1077.

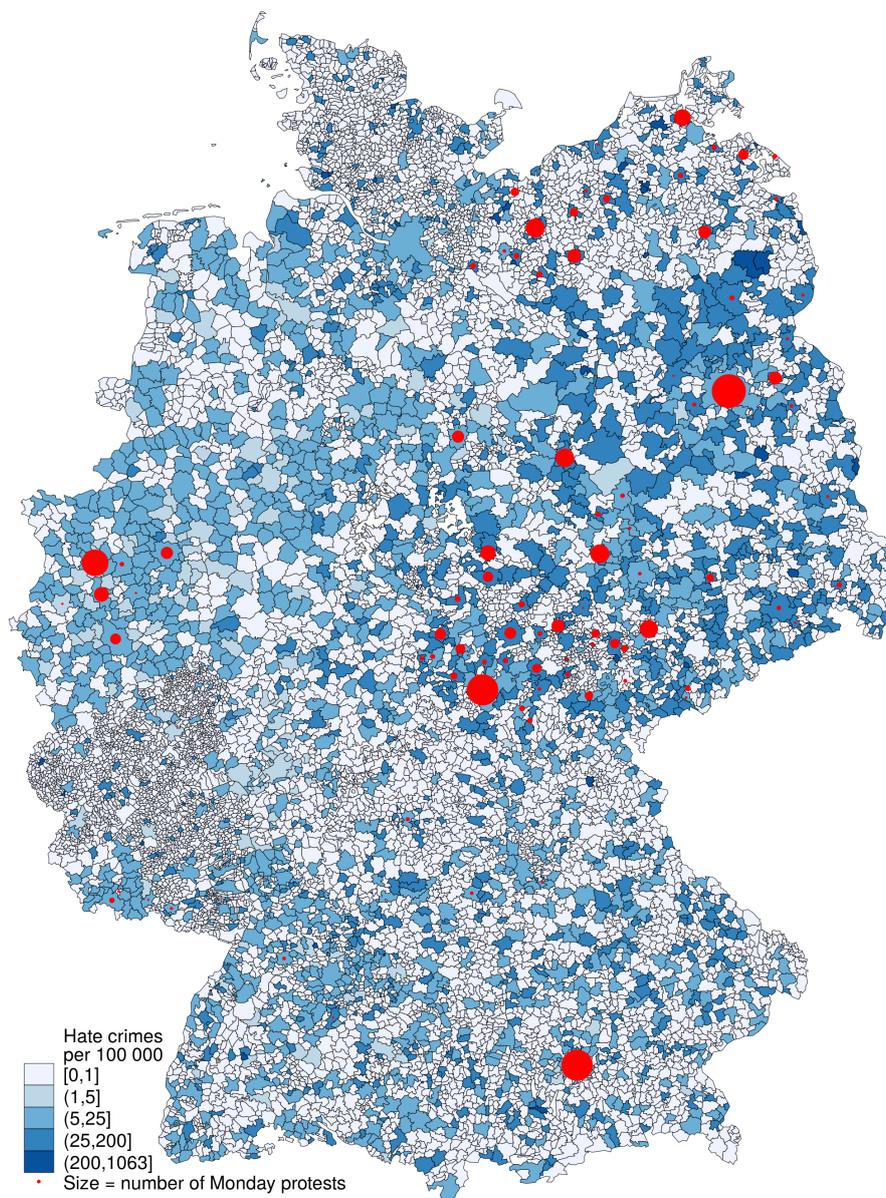
- Carr, J., Clifton-Sprigg, J., James, J., & Vujic, S. (2020). *Love Thy Neighbour? Brexit and Hate Crime*. IZA Discussion Paper DP No. 13902.
- Casanueva, A. (2021). Can chants in the street change parliament's tune? the effects of the 15M social movement on spanish elections. *mimeo*.
- Conley, T. G. (1999). Gmm estimation with cross sectional dependence. *Journal of Econometrics*, 92(1), 1–45.
- Connolly, K. (2014). Estimated 15,000 people join 'pinstriped nazis' on march in dresden. *The Guardian*, 15 Dec 2014.
- CSIS (2022). The evolution of domestic terrorism. Accessed: 2025-02-27.
- 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. (2023). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey. *The Econometrics Journal*, 26(3), C1–C30.
- Edwards, S. (2019). On latin american populism, and its echoes around the world. *Journal of Economic Perspectives*, 33(4), 76–99.
- Field, S. (1992). The effect of temperature on crime. *The British Journal of Criminology*, 32(3), 340–351.
- Flückiger, M. & Ludwig, M. (2025). The structure of online social networks and social movements: Evidence from the black lives matter protests. *Journal of Public Economics*, 246, 105373.
- 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.
- Fujiwara, T., Meng, K., & Vogl, T. (2016). Habit formation in voting: Evidence from rainy elections. *American Economic Journal: Applied Economics*, 8(4), 160–88.
- Funke, M., Schularick, M., & Trebesch, C. (2023). Populist leaders and the economy. *American Economic Review*, 113(12), 3249–3288.
- Gethin, A. & Pons, V. (2024). *Social Movements and Public Opinion in the United States*. Working Paper 32342, National Bureau of Economic Research.
- Gomez, B. T., Hansford, T. G., & Krause, G. A. (2007). The republicans should pray for rain: Weather, turnout, and voting in us presidential elections. *The Journal of Politics*, 69(3), 649–663.
- González, F. (2020). Collective action in networks: Evidence from the chilean student movement. *Journal of Public Economics*, 188, 104220.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
- Grabow, K. (2016). Pegida and the alternative für deutschland: two sides of the same coin? *European View*, 15, 173–181.
- Grosjean, P., Masera, F., & Yousaf, H. (2022). Inflammatory Political Campaigns and Racial Bias in Policing*. *The Quarterly Journal of Economics*, 138(1), 413–463.

- Guriev, S. & Papaioannou, E. (2022). The political economy of populism. *Journal of Economic Literature*, 60(3), 753–832.
- Haller, A. (2019). Paradoxical populism: how pegida relates to mainstream media. *Information, Communication & Society*.
- 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.
- 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).
- Jiménez Durán, R., Müller, K., & Schwarz, C. (2022). The Effect of Content Moderation on Online and Offline Hate: Evidence from Germany’s NetzDG. *Available at SSRN: <https://ssrn.com/abstract=4230296>*.
- Kanol, E. & Knoesel, J. (2021). Right-wing extremist mobilization in germany. *WZB Berlin Social Science Center, Version*, 1(0), 125–41.
- 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).
- Landesamt für Verfassungsschutz Sachsen (2021). Lfv stuft pegida als erwiesene extremistische bestrebung ein. Press release, 7 May 2021.
- Larreboure, M. & González, F. (2021). The impact of the women’s march on the us house election. *Pontificia Universidad Catolica de Chile,(Maret 2021)*, (pp. 4–5).
- Levitsky, S. & Ziblatt, D. (2019). *How democracies die*. Crown.
- Levy, R. & Mattsson, M. (2023). The effects of social movements: Evidence from# metoo. *Available at SSRN 3496903*.
- 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.
- 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.
- 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.

- Reuters (2015). Pegida’s anti-muslim calls shake up german politics. 14 Jan 2015.
- Reuters (2016a). German pegida founder in court over hate speech. 19 Apr 2016.
- Reuters (2016b). German voters batter merkel over migrant policy; afd gains. 13 Mar 2016.
- 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*.
- Sabet, N., Liebald, M., & Friebel, G. (2025). Terrorism and voting: The rise of right-wing populism in germany. *American Economic Journal: Economic Policy*, 17(3), 407–440.
- Sharoff, S. (2006). Open-source corpora: Using the net to fish for linguistic data. *International Journal of Corpus Linguistics*, 11(4), 435–462.
- START (2022). National consortium for the study of terrorism and responses to terrorism. global terrorism database 1970 - 2020 [data file]. <https://www.start.umd.edu/data-tools/GTD>.
- Time Magazine (2015). Meet the german activist leading the movement against ‘islamization’. Jan 2015 cover story.
- U.S. Department of State (2020). International religious freedom report - germany (2020).
- Virchow, F. (2016). Pegida: Understanding the emergence and essence of nativist protest in dresden. *Journal of Intercultural Studies*, 37(6), 541–555.
- Vorländer, H., Herold, M., & Schäller, S. (2018). *PEGIDA and new right-wing populism in Germany*. Springer.
- Washington Post (2016). A far-right party just won seats in three german state parliaments. 22 Mar 2016.
- Weisskircher, M. (2023). Protest and electoral breakthrough: Challenger party emergence in germany and austria. *German Politics*.
- 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.

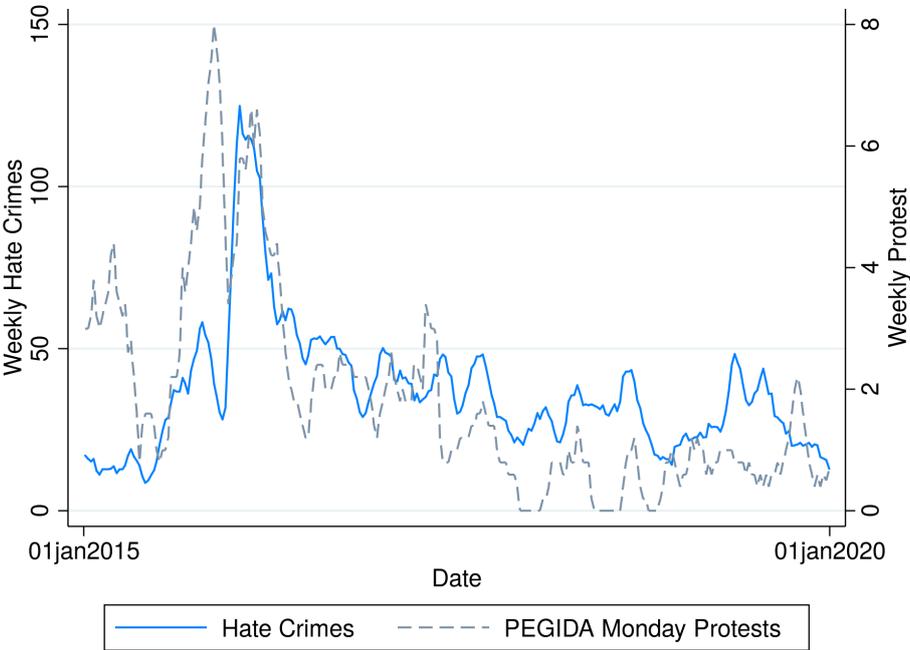
6 Tables and Figures

Figure 1. Hate crimes and PEGIDA protest across municipalities



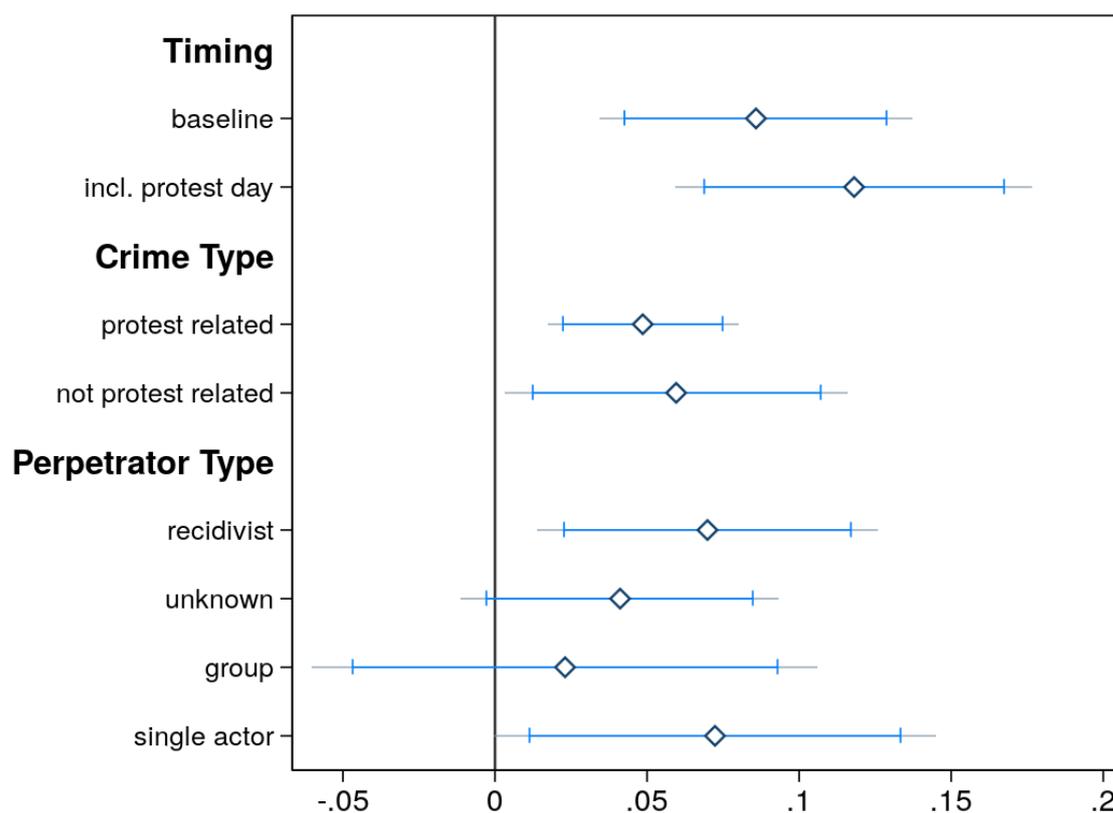
Notes: Map shows cumulative number of hate crimes (blue shades) and PEGIDA protests (red circles) across municipalities over the entire observation period (2015-December 2019).

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



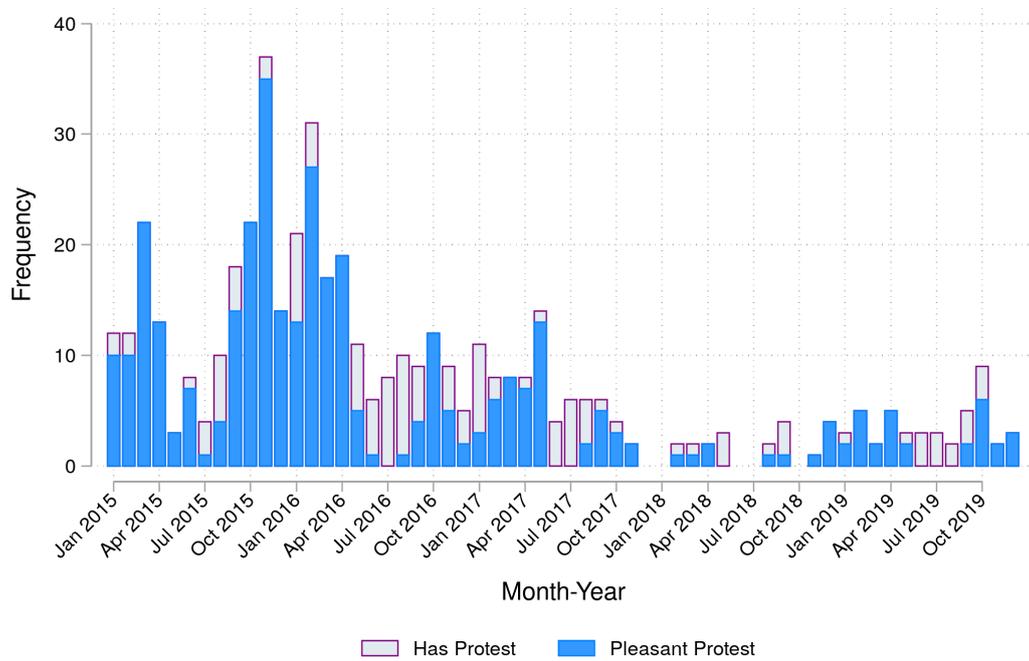
Notes: Figure shows the two-week moving average of the number of scheduled Monday PEGIDA protests and the number of hate crimes across all municipalities.

Figure 3. Effect of Protests on Hate-Crime Characteristics



Notes: Coefficient plot for reduced form regressions following our main estimating equation 1. Each line represents the estimate of a different regression for Protest \times Weather. Outcomes categorized by LLM based on descriptions of over 9,000 hate crimes. In order, each line represents: baseline specification for comparison; hate crimes of the same week excluding Monday; lines 3 and 4 identify whether the hate crimes was committed in a group or alone; lines 5 and 6 distinguish between hate crimes where the perpetrator was known to the police, has committed hate crimes before, or belonged to a known right-wing extremist group or not; line 7 and 8 examine whether the hate crime happened in a public space (like public squares and parks) or not and lines 9 and 10 whether the hate crime happened in the context of a protest or rally or not.

Figure 4. Time variation of PEGIDA protests on pleasant days



Notes: Graph shows the number of total Monday protests (light gray) and number of Monday protests on a pleasant day (blue) across all municipalities by month for the period January 2015 to December 2019.

Table 1. **Pleasant-Day Protest, Salience and Hate Crimes**

| | log turnout at Monday protest (1) | log Newspaper reports on protest (2) | log PEGIDA tweets (3) | Any hate crime Tue-Sun [0,1] (4) |
|------------------------|---|--|-----------------------------|--|
| Protest × Weather | 0.388*** (0.081) | 0.337*** (0.096) | 0.231** (0.116) | 0.086*** (0.026) |
| Protest | 4.393*** (0.128) | 0.354** (0.148) | 0.511*** (0.126) | 0.086** (0.043) |
| Weather | -0.000 (0.000) | 0.000 (0.000) | 0.003*** (0.000) | -0.00046*** (0.00015) |
| Observations | 2,820,944 | 2,820,944 | 2,762,218 | 2,820,944 |
| Municipalities | 10,825 | 10,825 | 10,599 | 10,825 |
| Municipality & Week FE | Yes | Yes | Yes | Yes |
| Protest × Controls | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes |
| Protest AME | 4.655*** (0.096) | 0.324*** (0.085) | 0.430*** (0.082) | 0.048** (0.021) |
| Mean dep. var. | 0.00076 | 0.000119 | 0.00830 | 0.00286 |

Notes: This table presents results from the main estimating equation 1, linking exogenous variation in pleasant weather on scheduled PEGIDA Monday protests to protest turnout, media attention, online engagement, and subsequent hate crimes. Column 1 reports the effect on the logarithm of one plus the Monday protest turnout; Columns 2 and 3 capture its impact on local and online salience, measured respectively by the log number of newspaper reports (from universe of newspaper articles from GENIOS) and PEGIDA-related tweets (from universe of geo-localized tweets containing the word PEGIDA). Column 4 displays the reduced-form relationship between pleasant weather on protest days and the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2. **Placebo and Validation Exercises**

| | presence of Monday protest | Any hate crime Tue-Sun [0,1] | | number of other protests | log. attendance other protests | |
|--------------------|-------------------------------|---------------------------------|--------------------------|-----------------------------|-----------------------------------|-----------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Protest × Weather | | | 0.086*** (0.026) | 0.078*** (0.027) | 0.091 (0.061) | -0.005 (0.092) |
| Protest | | | 0.086** (0.043) | 0.078* (0.042) | 0.101*** (0.039) | -0.130 (0.098) |
| Weather | -0.00002 (0.00002) | -0.00036*** (0.00012) | -0.00043*** (0.00016) | -0.00047*** (0.00015) | 0.00010 (0.00006) | -0.00002 (0.00006) |
| Observations | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 |
| Municipalities | 10,825 | 10,825 | 10,825 | 10,825 | 10,825 | 10,825 |
| Week FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Protest × Controls | | | Yes | Yes | Yes | Yes |
| Weather × Controls | | | Yes | Yes | Yes | Yes |
| Weather Tue–Sun | | | Yes | Yes | Yes | Yes |
| Protests Tue–Sun | | | | Yes | | Yes |
| Mean dep. var. | 0.000168 | 0.00286 | 0.000395 | 0.000395 | 0.000411 | 0.00168 |

Notes: This table reports placebo and validation exercises using alternative outcomes and specifications. Column 1 correlates the occurrence of scheduled Monday protests with pleasant weather. Columns 2 to 4 use an indicator for any hate crime between Tuesday and Sunday as the outcome. Column 2 correlates an indicator for any hate crime between Tuesday and Sunday with pleasant weather. Column 3 reports results from our main estimating equation 1, additionally controlling for weather conditions on other days of the week; Column 4 further controls for the number of unrelated right-wing protests occurring on other weekdays. Columns 5 and 6 use placebo outcomes: the number of unrelated right-wing protests on other weekdays (Column 5) and the logarithm of one plus total attendance at these protests (Column 6). These specifications additionally control for weather on other days of the week, and Column 6 also controls for the number of unrelated protests. All specifications include municipality and week fixed effects and, unless otherwise indicated, interactions of protest and weather with baseline controls. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3. **Diffusion: geographic and social media proximity to pleasant-day protest**

| | Any hate crime [0,1] | | | | |
|--------------------|----------------------------|------------------------------|------------------------|---------------------|----------------------------------|
| | geographic distance (1) | social media networks (2) | PEGIDA networks (3) | exposure all (4) | dep. var.: hate crime t+1 (5) |
| G_{it} | 0.009* (0.005) | | | 0.008 (0.005) | 0.001 (0.005) |
| S_{it}^W | | 0.042** (0.019) | | -0.006 (0.025) | -0.003 (0.028) |
| S_{it}^P | | | 0.087*** (0.018) | 0.090*** (0.021) | 0.113*** (0.022) |
| Observations | 2,820,469 | 2,820,469 | 2,820,469 | 2,820,469 | 2,809,645 |
| Municipalities | 10,825 | 10,825 | 10,825 | 10,825 | 10,825 |
| Week FE | Yes | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.00279 | 0.00279 | 0.00279 | 0.00279 | 0.00280 |

Notes: This table presents results from our spatial diffusion estimating equation 2. Sample includes all municipality-week units without a PEGIDA protest on the Monday of that week. G_{it} is the geographic proximity weighted exposure to pleasant-day protest in other municipalities (standardized). S_{it} is the social media network the pleasant-day exposure to pleasant-day protests in other municipalities (standardized). Superscript P indicates PEGIDA network, superscript W indicates wider network, based on random sample of tweets. Details on the construction can be found in Section 4.1 and Appendix C. Outcome in columns 1 to 4 is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week, using data from ProAsyl and AAF. Outcome in column 5 reports any hate crime in the following week. Coefficients multiplied by 100 for presentation. Time horizon is January 2015 until December 2019. Observations are municipality-week units. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01.

Table 4. **Counter-Protests, Offline and Online**

| | counterprotest [0,1] | | log(tweets refugees welcome) | |
|---------------------------|----------------------|--------------------|------------------------------|--------------------|
| | t (1) | t+1 (2) | t (3) | t+1 (4) |
| Protest \times Weather | 0.117*** (0.034) | 0.037 (0.036) | 0.279** (0.124) | 0.188 (0.163) |
| Observations | 2,820,944 | 2,810,119 | 2,258,048 | 2,247,235 |
| Municipalities | 10,825 | 10,825 | 10,813 | 10,813 |
| Week FE | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes |
| Protest \times Controls | Yes | Yes | Yes | Yes |
| Weather \times Controls | Yes | Yes | Yes | Yes |
| Protest AME | 0.136*** (0.029) | 0.059** (0.024) | 0.179*** (0.066) | 0.154** (0.078) |
| Mean dep. var. | 0.000509 | 0.000499 | 0.0159 | 0.0160 |

Notes: This table presents results from our main estimating equation 1, linking exogenous variation in pleasant weather on scheduled PEGIDA Monday protests to protest turnout, media attention, online engagement, and subsequent hate crimes. Outcomes in columns 1 and 2 capture whether a counter-protest occurred in the same week and the following week, using data from the newspaper archive GENIOS. Columns 3 and 4 measure the log number of geo-localized #refugeeswelcome tweets in the same week and the subsequent week. The table also reports the average marginal effect (AME) of a protest evaluated at the mean of the controls, with standard errors in parentheses. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5. **Exposure to pro- and anti-PEGIDA coverage**

| | Any hate crime [0,1] | | | |
|------------------------------------|----------------------|---------------------|---------------------|---------------------|
| | t (1) | t (2) | t (3) | t+1 (4) |
| Twitter Exposure ^{pro} | 0.053*** (0.013) | | 0.053*** (0.013) | 0.059*** (0.018) |
| Twitter Exposure ^{anti} | 0.005 (0.033) | | 0.005 (0.033) | -0.006 (0.036) |
| Newspaper Exposure ^{pro} | | 0.042*** (0.013) | 0.041*** (0.013) | 0.042*** (0.011) |
| Newspaper Exposure ^{anti} | | -0.016* (0.009) | -0.018* (0.009) | -0.021** (0.009) |
| Observations | 2,257,618 | 2,809,456 | 2,248,737 | 2,248,737 |
| Municipalities | 10,813 | 10,784 | 10,772 | 10,772 |
| Week FE | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes |
| Protest × Controls | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.00293 | 0.00280 | 0.00294 | 0.00294 |

Notes: This table presents results from our media exposure estimating equation 3. Sample includes all municipality-week units without a PEGIDA protest on the Monday of that week. Column 1 includes exposure to pro and anti-PEGIDA messaging on Twitter, and column 2 includes exposure to newspaper articles about pro and anti-PEGIDA protests, as defined in section 4.2. Outcome in columns 1 to 3 is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week, using data from ProAsyl and AAF. Outcome in column 4 reports any hate crime in the following week. All baseline controls of main specification 1 are included. Coefficients multiplied by 100 for presentation. Time horizon is January 2015 until December 2019. Observations are municipality-week units. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01.

Online Appendix

Appendix A: Additional Results

A.1 Protest at the extensive margin

Our main research design identifies the effect of protest *saliency*, which is plausibly the policy-relevant dimension since marginal participants and journalists may be more likely to respond to information about the potential negative externalities of attendance and coverage. In this section, we offer a complementary approach that studies the extensive margin: whether the *occurrence* of an anti-immigration protest increases subsequent hate crimes. This may be the relevant estimand if policy-makers consider banning some forms of protest altogether.

As mentioned above, the German Assembly Act creates a challenge for studying protest at the extensive margin since short-run shocks, including weather conditions, do not predict the occurrence of protest. This institutional feature motivates a few adjustments relative to our previous specification. We need to identify an unexpected and large enough protest trigger that shifts the demand for PEGIDA protests across the entire country and the time window that allows for the organization and registration of a PEGIDA Monday protest.

To obtain plausibly exogenous variation in the demand for PEGIDA protest, we exploit the timing of Islamist terrorist attacks in the rest of Europe as national-level shocks that can trigger anti-immigration mobilization in Germany. PEGIDA organizers explicitly referenced prominent attacks, such as the *Charlie Hebdo* shooting, as motivations for mobilization, suggesting that the extensive margin of protest responds to such events even when they occur outside Germany. We use the Global Terrorism Database (START, 2022) to identify Islamist terrorist attacks between 2015 and 2019. Table C.3 reports the date and location of these events. We exclude attacks occurring within Germany to avoid the local variation in hate-crimes and anti-minority sentiment triggered by Islamic terrorist attacks within Germany.

Identification comes from the interaction between this unexpected protest trigger and heterogeneity in municipalities' propensity to host PEGIDA protests. We therefore construct a PEGIDA propensity score, based on baseline municipal characteristics. Concretely, we estimate a predictive LASSO model for PEGIDA protest activity that predicts the number of PEGIDA protests between 2015 and 2019 *in normal times* i.e., protests that do not happen in the aftermath of an Islamist terror attack.²⁹ We standardize the predicted propensity for ease of interpretation. In addition, to ensure that we are not simply comparing municipalities that are differentially prone to hate crimes, we include the interaction between the terrorist-attack indicator with the overall crime rate at baseline.³⁰ This allows the effect to vary flexibly with baseline crime prevalence, so that the remaining comparison isolates differential responses in PEGIDA-prone versus non-prone municipalities rather than in crime-prone versus non-prone municipalities. We estimate the following first-stage and second-stage specification at the municipality and week level:

$$\text{Protest}_{i,t} = \alpha_i + \lambda_t + \beta \text{Attack}_t \times \text{Propensity}_i + \kappa \text{Attack}_t \times \bar{C}_{i,t_0} + \theta \text{Weather}_{i,t} + \tilde{\pi}'_{it} L_{it} + \epsilon_{it}, \quad (4)$$

$$H_{i,t} = \omega_i + \tau_t + \gamma \widehat{\text{Protest}}_{i,t} + \phi \text{Attack}_t \times \bar{C}_{i,t_0} + \chi \text{Weather}_{i,t} + \tilde{\pi}'_{it} L_{it} + \varepsilon_{it}, \quad (5)$$

where Protest_{it} equals to one if there was a PEGIDA Monday protest in municipality i at time t and Attack_t equals to one if there was an Islamic Terrorist Attack during the four week window from $t - 4$

²⁹This ensures that the propensity score reflects baseline correlates of PEGIDA mobilization rather than protest activity mechanically induced by terrorist attacks. We use the fitted values from this model, which captures how PEGIDA-prone a municipality is given its baseline covariates.

³⁰Since information on hate crimes only starts in 2015, we use the overall crime rate as a proxy.

to $t - 1$. $H_{i,t}$ equals one if at least one hate crime against refugees is recorded in municipality i during the six days (Tuesday to Sunday) following date t (excluding the Monday of the potential protest) and zero otherwise. \bar{C}_{i,t_0} is the overall crime rate at baseline, also interacted with the indicator for Islamic Terrorist Attack. We include municipality fixed effects and week fixed effects; the latter absorb the aggregate effect of terrorist-attack news common to all municipalities on a given date. Analogous to our previous specification, we control for local weather conditions on the protest day, as well as for a vector of control $\tilde{\pi}'_{it}L_{it}$ which accounts for lagged weather, protest, and pleasant-day protest, along with their interactions with contemporaneous protest and weather conditions. Standard errors are clustered by municipality.

Table A.1 presents our results. Column 1 shows that the interaction predicts the occurrence of PEGIDA protest. The probability of observing a PEGIDA protest within 28 days after a Islamic terrorist attack increases by 0.16 percentage points for each standard deviation in PEGIDA propensity, with a Kleinbergen-Paap F-Statistic of 70.77. Column 2 reports the OLS correlation between protest and hate crimes, showing that the probability of recording a hate crime increases by 7.8 percentage points in the aftermath of a protest. Columns 3 and 4 report the 2SLS and reduced form coefficients. The former suggests that the occurrence of a PEGIDA protest increases the probability of a hate crime on the following six days by 90 p.p.

The exclusion restriction is violated if the interaction of propensity to protest and terrorist attacks causes hate crimes in a way that is not mediated by the presence of protests. There are two ways this can happen. First, hate crimes could be caused directly: for example, the propensity to protest may be linked to the propensity to commit hate crimes (that is not captured by the overall crime rate at baseline) and radicalized individuals may be more likely to commit hate crimes in the aftermath of a terrorist attack. To address this possibility, we consider hate crimes that occur after a terrorist attack but before a PEGIDA Monday protest. Hence, in column 5, we omit all municipality-week observations with protests. Reassuringly, we do not see a surge in hate crimes in the weeks following the terrorist attacks without corresponding protests. The effect sets in only after the protest, suggesting that we capture the direct effect of protest rather than differential responses to Islamic terrorist attacks.

Table A.1. **Extensive margin of protest**

| | Protest | Any hate crime [0,1] | | | |
|---------------------|-------------------------|----------------------|---------------------|-------------------------|----------------------|
| | | all | all | all | before protest |
| | | First Stage | OLS | 2SLS | Reduced Form |
| | (1) | (2) | (3) | (4) | (5) |
| Attack × Propensity | 0.00158*** (0.00019) | | | 0.00142*** (0.00034) | 0.00052 (0.00064) |
| Protest | | 0.078*** (0.025) | 0.898*** (0.264) | | |
| Observations | 2,810,243 | 2,810,243 | 2,810,243 | 2,810,243 | 2,809,768 |
| Municipalities | 10,784 | 10,784 | 10,784 | 10,784 | 10,784 |
| Week FE | Yes | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes | Yes |
| Weather | Yes | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.000169 | 0.00286 | 0.00286 | 0.00286 | 0.00280 |
| KPF stat. | 70.77 | | | | |

Notes: Estimate of the extensive margin of protest, instrumented by a dummy variable indicating whether an Islamist terrorist attack took place in Europe in the previous 4 weeks and the propensity of a municipality to protest, estimated from LASSO-selected baseline characteristics. Column 1 presents the first stage estimate, with presence of a Monday protest as outcome. Column 2 shows the OLS effect of protest on hate crime, and column 3 the 2SLS estimate. Column 4 shows the reduced form effect of the instrument on hate crime. Column 5, as a placebo exercise, shows the effect of the instrument of hate crime, removing all observations after protest. All columns include week and municipality fixed effects, and control by the weather on the day of the protest, lagged variables, and their interaction. Time horizon is January 2015 until December 2019. Observations are municipality-week units. SE clustered by municipality ; *p<0.1; **p<0.05; ***p<0.01. Treatment is measured as the presence of 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 in the week, excluding crimes committed on Monday. Kleinbergen-Paap F-Statistic is reported at the bottom of the table for column 1.

A.2 Other results

Table A.2. Descriptive statistics

| | mean | sd | min | max |
|---|----------|-----------|----------|-----------|
| Main variables: | | | | |
| Any HC | 0.0029 | 0.0534 | 0.0000 | 1.0000 |
| Protest | 0.0002 | 0.0130 | 0.0000 | 1.0000 |
| Total Participants | 0.0231 | 2.8829 | 0.0000 | 2300 |
| Total Participants — Protest | 137.4211 | 174.7641 | 4.0000 | 2300 |
| Log(participants) | 0.0008 | 0.0598 | 0.0000 | 7.7411 |
| Log Participants — Protest | 4.5358 | 0.8198 | 1.6094 | 7.7411 |
| Weather | 0.6772 | 0.4676 | 0.0000 | 1.0000 |
| Protest × Weather | 0.0001 | 0.0111 | 0.0000 | 1.0000 |
| Terrorist attack last 4 weeks | 0.6130 | 0.4871 | 0.0000 | 1.0000 |
| Reactivity to attacks | 0.0000 | 1.0000 | -3.0998 | 100.5166 |
| Types of hate crimes: | | | | |
| incl. Monday | 0.0033 | 0.0574 | 0.0000 | 1.0000 |
| protest related | 0.0001 | 0.0103 | 0.0000 | 1.0000 |
| protest unrelated | 0.0028 | 0.0526 | 0.0000 | 1.0000 |
| known perp. | 0.0026 | 0.0512 | 0.0000 | 1.0000 |
| unknown perp. | 0.0003 | 0.0168 | 0.0000 | 1.0000 |
| group perp. | 0.0012 | 0.0349 | 0.0000 | 1.0000 |
| single perp. | 0.0010 | 0.0314 | 0.0000 | 1.0000 |
| public space | 0.0004 | 0.0195 | 0.0000 | 1.0000 |
| no public space | 0.0025 | 0.0503 | 0.0000 | 1.0000 |
| Media and Social media measures: | | | | |
| log Newspaper reports | 0.0007 | 0.0392 | 0.0000 | 6.6657 |
| log Pegida tweets | 0.0389 | 0.2023 | 0.0000 | 8.9804 |
| Counterprotest Presence | 0.0005 | 0.0226 | 0.0000 | 1.0000 |
| log Counterprotest attendance | 0.0019 | 0.1149 | 0.0000 | 13.1343 |
| log Tweets #RefugeesWelcome | 0.0159 | 0.1308 | 0.0000 | 8.2169 |
| S_{it}^P | 0.9737 | 5.0790 | 0.0000 | 902.7473 |
| S_{it}^W | 4.88e+05 | 1.93e+06 | -0.0000 | 2.87e+08 |
| G_{it} | 264.4453 | 1239.3387 | 0.0000 | 3.30e+04 |
| Twitter Exposure ^{pro} | 0.0000 | 1.0000 | -0.0221 | 629.7974 |
| Twitter Exposure ^{anti} | 0.0000 | 1.0000 | -0.0171 | 774.1576 |
| Newspaper Exposure ^{pro} | 0.0000 | 1.0000 | -0.2024 | 55.5187 |
| Newspaper Exposure ^{anti} | 0.0000 | 1.0000 | -0.1836 | 60.5411 |
| Additional controls: | | | | |
| GDP per capita | 3.03e+04 | 8633.9761 | 1.51e+04 | 1.88e+05 |
| Population density | 185.9822 | 284.2154 | 0.0000 | 4736.1055 |
| Unemployment rate | 0.0211 | 0.0130 | 0.0000 | 0.2903 |
| Refugees share | 0.0123 | 0.0061 | 0.0008 | 0.1300 |
| AFD vote share | 0.0998 | 0.0623 | 0.0222 | 0.3502 |
| Crime rate (per 100 000) | 11.0900 | 24.5935 | 0.0371 | 1182.4286 |
| Observations | 2820944 | | | |

Notes: Descriptive statistics for the main variables used in the analysis for all German municipalities averaged over the whole period of analysis. Variables Twitter Exposure^{pro}, Twitter Exposure^{anti}, Newspaper Exposure^{pro}, Newspaper Exposure^{anti}, S_{it}^P , S_{it}^W and G_{it} are standardized. The definition of the last three variables can be found in the notes of Table 3.

Table A.3. Descriptive statistics by municipality type

| | Municipalities with at least one protest | | | | Municipalities with no protest | | | |
|---|--|-----------|----------|-----------|--------------------------------|-----------|----------|-----------|
| | mean | sd | min | max | mean | sd | min | max |
| Main variables: | | | | | | | | |
| Any HC | 0.0827 | 0.2754 | 0.0000 | 1.0000 | 0.0022 | 0.0470 | 0.0000 | 1.0000 |
| Protest | 0.0212 | 0.1441 | 0.0000 | 1.0000 | 0.0000 | 0.0000 | 0.0000 | 0.0000 |
| Total Participants | 2.9150 | 32.2278 | 0.0000 | 2300.0000 | 0.0000 | 0.0000 | 0.0000 | 0.0000 |
| Total Participants — Protest | 137.4211 | 174.7641 | 4.0000 | 2300.0000 | . | . | . | . |
| Log(participants) | 0.0962 | 0.6644 | 0.0000 | 7.7411 | 0.0000 | 0.0000 | 0.0000 | 0.0000 |
| Log Participants — Protest | 4.5358 | 0.8198 | 1.6094 | 7.7411 | . | . | . | . |
| Weather | 0.6739 | 0.4688 | 0.0000 | 1.0000 | 0.6772 | 0.4675 | 0.0000 | 1.0000 |
| Protest × Weather | 0.0156 | 0.1240 | 0.0000 | 1.0000 | 0.0000 | 0.0000 | 0.0000 | 0.0000 |
| Terrorist attack last 4 weeks | 0.6133 | 0.4870 | 0.0000 | 1.0000 | 0.6130 | 0.4871 | 0.0000 | 1.0000 |
| Reactivity to attacks | 2.8355 | 10.4790 | -0.3877 | 100.5166 | -0.0227 | 0.2540 | -3.0998 | 10.9327 |
| Types of hate crimes: | | | | | | | | |
| incl. Monday | 0.0933 | 0.2908 | 0.0000 | 1.0000 | 0.0026 | 0.0507 | 0.0000 | 1.0000 |
| protest related | 0.0056 | 0.0745 | 0.0000 | 1.0000 | 0.0001 | 0.0079 | 0.0000 | 1.0000 |
| protest unrelated | 0.0795 | 0.2705 | 0.0000 | 1.0000 | 0.0022 | 0.0464 | 0.0000 | 1.0000 |
| known perp. | 0.0779 | 0.2681 | 0.0000 | 1.0000 | 0.0020 | 0.0449 | 0.0000 | 1.0000 |
| unknown perp. | 0.0089 | 0.0939 | 0.0000 | 1.0000 | 0.0002 | 0.0147 | 0.0000 | 1.0000 |
| group perp. | 0.0394 | 0.1945 | 0.0000 | 1.0000 | 0.0009 | 0.0302 | 0.0000 | 1.0000 |
| single perp. | 0.0341 | 0.1814 | 0.0000 | 1.0000 | 0.0007 | 0.0269 | 0.0000 | 1.0000 |
| public space | 0.0161 | 0.1259 | 0.0000 | 1.0000 | 0.0003 | 0.0159 | 0.0000 | 1.0000 |
| no public space | 0.0715 | 0.2577 | 0.0000 | 1.0000 | 0.0020 | 0.0445 | 0.0000 | 1.0000 |
| Media and Social media measures: | | | | | | | | |
| log Newspaper reports | 0.0559 | 0.3730 | 0.0000 | 6.6657 | 0.0003 | 0.0203 | 0.0000 | 4.6540 |
| log Pegida tweets | 0.6088 | 1.2175 | 0.0000 | 8.9804 | 0.0343 | 0.1627 | 0.0000 | 7.8184 |
| Counterprotest Presence | 0.0279 | 0.1647 | 0.0000 | 1.0000 | 0.0003 | 0.0170 | 0.0000 | 1.0000 |
| log Counterprotest attendance | 0.1260 | 0.9463 | 0.0000 | 13.1343 | 0.0009 | 0.0776 | 0.0000 | 10.5967 |
| log Tweets #RefugeesWelcome | 0.3476 | 0.8644 | 0.0000 | 8.2169 | 0.0133 | 0.1019 | 0.0000 | 7.2101 |
| S_{it}^P | 4.8347 | 30.9337 | 0.0000 | 902.7473 | 0.9428 | 4.2692 | 0.0000 | 756.0439 |
| S_{it}^W | 1.70e+06 | 8.52e+06 | -0.0000 | 2.87e+08 | 4.78e+05 | 1.78e+06 | -0.0000 | 1.75e+08 |
| G_{it} | 664.1487 | 2033.4717 | 0.0000 | 2.90e+04 | 261.2470 | 1230.3970 | 0.0000 | 3.30e+04 |
| Twitter Exposure ^{pro} | 1.0627 | 9.9457 | -0.0220 | 629.7974 | -0.0085 | 0.4554 | -0.0221 | 294.2334 |
| Twitter Exposure ^{anti} | 0.8331 | 10.0214 | -0.0171 | 774.1576 | -0.0067 | 0.4459 | -0.0171 | 361.6820 |
| Newspaper Exposure ^{pro} | 0.1339 | 1.2599 | -0.2024 | 31.2954 | -0.0011 | 0.9976 | -0.2024 | 55.5187 |
| Newspaper Exposure ^{anti} | 0.1025 | 1.2297 | -0.1836 | 37.0288 | -0.0008 | 0.9979 | -0.1836 | 60.5411 |
| Additional controls: | | | | | | | | |
| GDP per capita | 3.16e+04 | 1.23e+04 | 1.92e+04 | 8.55e+04 | 3.03e+04 | 8597.6489 | 1.51e+04 | 1.88e+05 |
| Population density | 739.4235 | 845.3701 | 41.6208 | 4736.1055 | 181.5538 | 270.6218 | 0.0000 | 4118.9292 |
| Unemployment rate | 0.0440 | 0.0147 | 0.0124 | 0.1003 | 0.0209 | 0.0128 | 0.0000 | 0.2903 |
| Refugees share | 0.0142 | 0.0091 | 0.0008 | 0.0493 | 0.0123 | 0.0060 | 0.0008 | 0.1300 |
| AFD vote share | 0.1395 | 0.0825 | 0.0352 | 0.3502 | 0.0994 | 0.0620 | 0.0222 | 0.3502 |
| Crime rate (per 100 000) | 0.6604 | 1.6928 | 0.0613 | 16.2552 | 11.1738 | 24.6738 | 0.0371 | 1182.4286 |
| Observations | 22393 | | | | 2798551 | | | |

Notes: Descriptive statistics for the main variables used in the analysis for all German municipalities averaged over the whole period of analysis, and split between municipalities with no or at least one protest. Variables Twitter Exposure^{pro}, Twitter Exposure^{anti}, Newspaper Exposure^{pro}, Newspaper Exposure^{anti}, S_{it}^P , S_{it}^W and G_{it} are standardized. The definition of the last three variables can be found in the notes of Table 3.

Table A.4. **Pleasant-Day Protest, Salience and Hate Crimes: Ever-treated only**

| | log turnout at Monday protest (1) | log Newspaper reports on protest (2) | log PEGIDA tweets (3) | Any hate crime Tue-Sun [0,1] (4) |
|---------------------------|---|--|-----------------------------|--|
| Protest \times Weather | 0.38623*** (0.07969) | 0.313*** (0.092) | 0.142 (0.102) | 0.08318*** (0.02631) |
| Protest | 4.38960*** (0.12744) | 0.347** (0.139) | 0.417*** (0.120) | 0.07015 (0.04469) |
| Weather | -0.00493* (0.00258) | -0.003 (0.010) | 0.024 (0.016) | -0.00930 (0.00867) |
| Observations | 22,393 | 22,393 | 22,393 | 22,393 |
| Municipalities | 86 | 86 | 86 | 86 |
| Municipality & Week FE | Yes | Yes | Yes | Yes |
| Protest \times Controls | Yes | Yes | Yes | Yes |
| Weather \times Controls | Yes | Yes | Yes | Yes |
| Protest AME | 4.640*** (0.097) | 0.283*** (0.073) | 0.273*** (0.062) | 0.044** (0.022) |
| Mean dep. var. | 0.0962 | 0.00168 | 0.0211 | 0.0827 |

Notes: This table presents results from the main estimating equation 1, linking exogenous variation in pleasant weather on scheduled PEGIDA Monday protests to protest turnout, media attention, online engagement, and subsequent hate crimes. The sample is limited to municipalities that experienced at least one PEGIDA Monday protest over the study period. Column 1 reports the effect on the logarithm of one plus the Monday protest turnout; Columns 2 and 3 capture its impact on local and online salience, measured respectively by the log number of newspaper reports (from universe of newspaper articles from GENIOS) and PEGIDA-related tweets (from universe of geo-localized tweets containing the word PEGIDA). Column 4 displays the reduced-form relationship between pleasant weather on protest days and the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5. **2SLS and OLS estimates of the effect of protest salience on hate crimes**

| | Any hate crime Tue-Sun [0,1] | | | |
|-----------------------|------------------------------|----------------------|----------------------|----------------------|
| | 2SLS | | OLS | |
| | (1) | (2) | (3) | (4) |
| Log(participants) | 0.221** (0.088) | | -0.021 (0.022) | |
| PC1: Protest Salience | | 0.011** (0.005) | | 0.002*** (0.001) |
| Protest | -0.885** (0.396) | -0.313* (0.163) | 0.224* (0.124) | 0.049 (0.050) |
| Weather | -0.000*** (0.000) | -0.001*** (0.000) | -0.000*** (0.000) | -0.000*** (0.000) |
| Observations | 2,820,944 | 2,762,218 | 2,820,944 | 2,762,218 |
| Municipalities | 10825 | 10599 | 10825 | 10599 |
| Week FE | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes |
| Protest × Controls | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes |
| First stage coeff. | 0.388*** (0.081) | 7.463*** (1.631) | | |
| Mean dep. var. | 0.00286 | 0.00291 | 0.00286 | 0.00291 |
| KPF-first | 23.11 | 20.93 | | |

Note: 2SLS (columns 1 and 2) and OLS (columns 3 and 4) estimation of the effect of attendance in protests and protest salience on hate crimes. Instrument is pleasant weather dummy as defined in section 3.1 interacted with scheduled Monday protest, controlling for protest and weather separately. All columns include week and municipality fixed effects. Column 1 instruments the logarithm of one plus the number of participants, column 2 instruments standardized protest salience, defined as the first principal component of the logarithm of protest attendance, of newspaper reports, and of PEGIDA tweets. Column 3 and 4 show the equivalent OLS results. Time horizon is January 2015 until December 2019. Observations are municipality-week units. SE clustered by municipality; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. 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 column 4 Kleinbergen-Paap F-Statistics and first stage coefficients are reported at the bottom of the table. Controls comprise the lagged pleasant weather, lagged protest, and lagged pleasant-day protest as well as the interaction between those controls and the contemporaneous weather and protest dummy.

Table A.6. **Diffusion: All Municipalities**

| | Any hate crime [0,1] | | | | |
|---------------------------|-------------------------------|---------------------------------|---------------------------|------------------------|-------------------------------------|
| | geographic distance (1) | social media networks (2) | PEGIDA networks (3) | exposure all (4) | dep. var.: hate crime t+1 (5) |
| G_{it} | 0.009* (0.005) | | | 0.009 (0.005) | 0.001 (0.005) |
| S_{it}^W | | 0.048** (0.020) | | -0.005 (0.025) | -0.004 (0.028) |
| S_{it}^P | | | 0.095*** (0.021) | 0.097*** (0.023) | 0.124*** (0.023) |
| Protest \times Weather | 0.086*** (0.026) | 0.084*** (0.026) | 0.081*** (0.026) | 0.081*** (0.026) | 0.032 (0.051) |
| Observations | 2,820,469 | 2,820,469 | 2,820,469 | 2,820,469 | 2,809,645 |
| Municipalities | 10825 | 10825 | 10825 | 10825 | 10825 |
| Week FE | Yes | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes | Yes |
| Protest \times Controls | Yes | Yes | Yes | Yes | Yes |
| Weather \times Controls | Yes | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.00286 | 0.00286 | 0.00286 | 0.00286 | 0.00286 |

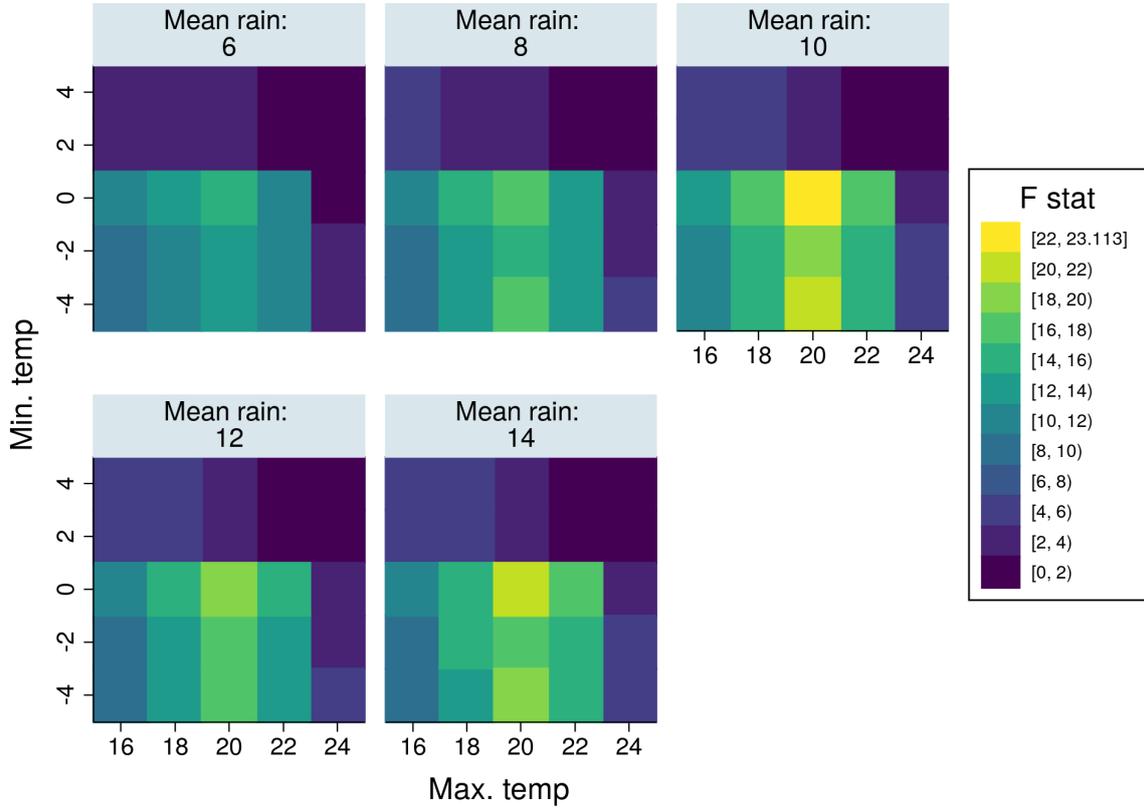
Notes: This table presents results from our spatial diffusion estimating equation 2, replicating Table 3 but including all municipality-week observations. G_{it} is the geographic proximity weighted exposure to pleasant-day protest in other municipalities (standardized). S_{it} is the social media network the pleasant-day exposure to pleasant-day protests in other municipalities (standardized). Superscript P indicates PEGIDA network, superscript W indicates wider network, based on random sample of tweets. Details on the construction can be found in Section 4.1 and Appendix C. Outcome in columns 1 to 4 is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week, using data from ProAsyl and AAF. Outcome in column 5 reports any hate crime in the following week. Coefficients multiplied by 100 for presentation. Time horizon is January 2015 until December 2019. Observations are municipality-week units. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01.

Table A.7. **Exposure to pro- and anti-PEGIDA coverage: All Municipalities**

| | Any hate crime [0,1] | | | |
|------------------------------------|----------------------|---------------------|---------------------|---------------------|
| | t (1) | t (2) | t (3) | t+1 (4) |
| Twitter Exposure ^{pro} | 0.035** (0.013) | | 0.035** (0.013) | 0.043** (0.018) |
| Twitter Exposure ^{anti} | -0.001 (0.025) | | -0.001 (0.025) | -0.008 (0.028) |
| Newspaper Exposure ^{pro} | | 0.043*** (0.013) | 0.042*** (0.013) | 0.042*** (0.012) |
| Newspaper Exposure ^{anti} | | -0.018* (0.009) | -0.020** (0.010) | -0.021** (0.009) |
| Protest × Weather | 0.067** (0.030) | 0.086*** (0.026) | 0.067** (0.030) | 0.081* (0.043) |
| Observations | 2,258,048 | 2,809,931 | 2,249,167 | 2,249,167 |
| Municipalities | 10,813 | 10,784 | 10,772 | 10,772 |
| Week FE | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes |
| Protest × Controls | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.00300 | 0.00286 | 0.00301 | 0.00301 |

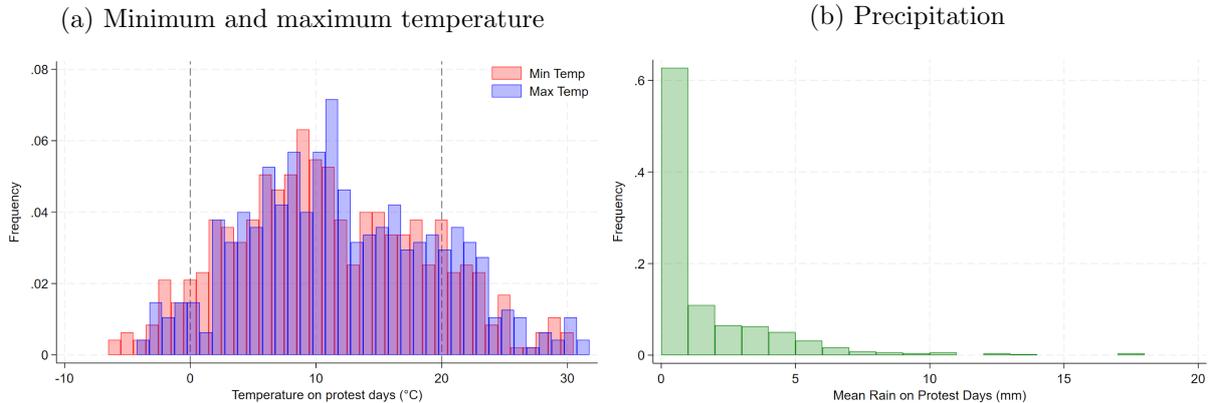
Notes: This table presents results from our media exposure estimating equation 3, replicating Table 5 but including all municipality-week observations. Column 1 includes exposure to pro and anti-PEGIDA messaging on Twitter, and column 2 includes exposure to newspaper articles about pro and anti-PEGIDA protests, as defined in section 4.2. Outcome in columns 1 to 3 is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week, using data from ProAsyl and AAF. Outcome in column 4 reports any hate crime in the following week. All baseline controls of main specification 1 are included. Coefficients multiplied by 100 for presentation. Time horizon is January 2015 until December 2019. Observations are municipality-week units. SE clustered by municipality; *p<0.1; **p<0.05; ***p<0.01.

Figure A.1. Variations on the weather instrument



Notes: Heatmap of Kleinbergen-Paap F-statistics from the first stage regression $\log(1 + Participants_{it}) = \alpha_i + \lambda_t + \gamma_1 (Weather_{it} \times Protest_{it}) + \gamma_2 Protest_{it} + \gamma_3 Weather_{it} + \tilde{\pi}'_{it} L_{it} + \nu_{it}$, where the outcome is the logarithm of one plus the number of participants at a given Monday PEGIDA protest. Y-axis reflect minimum temperature cut-offs from -4 to +4 Celcius. X-axis reflects maximum temperature cut-offs ranging from 16 to 24 Celcius and boxes vary the rain cut-off between 6 and 14 mm per square meter. All are measured between 2pm and 6pm on the protest day, which corresponds to the hours leading up to and the beginning of the protest, which typically start between 4:30 and 6 pm.

Figure A.2. Weather on protest days



Note: Distribution of minimum temperature, maximum temperature (in degrees Celsius) and mean precipitation per hour (in mm) between 2pm and 6pm on protest days.

Appendix B: Robustness Checks

Serial correlation, intertemporal substitution and past protests. Table B.1 reports robustness checks that rule out alternative explanations related to serial past protest history and seasonal patterns. Across all specifications, standard errors are clustered at the municipality level and we continue to include week and municipality fixed effects plus the usual controls interacted with both protest and weather. Columns 1 and 2 include, respectively, the cumulative count of past PEGIDA Monday protests and the cumulative sum of past Monday protest turnout in each municipality; neither historical variable attenuates our main interaction term. This suggests that we are not merely capturing serially correlated past shocks and that conditional on past protest activity any additional pleasant-day protest generates a surge in local hate crimes of similar magnitude. Column 3 introduces municipality-by-month-of-year fixed effects to absorb any location-specific seasonal cycles. The interaction coefficient increases slightly, confirming that municipality-level seasonality does not drive our findings. It is also worth noting that the inclusion of this fixed effect accounts for the average pleasant weather probability in a given municipality and month. Thus, the pleasant weather dummy would capture a deviation from this average probability and thus captures a weather shock.

Event Study. Figure B.1 plots dynamic estimates of the effect of pleasant-day PEGIDA protests on the weekly probability of a hate crime in the same municipality, spanning 6 weeks before to 8 weeks after each protest. The point estimates (dots) at negative leads (weeks -6 to -1) lie tightly around zero with tight 90% and 95% confidence bands (dark and light vertical lines), indicating no pre-trend or Ashenfelter “dip” in hate-crime incidence prior to treatment. Starting in the week of the pleasant-day protest (week 0) and extending through subsequent weeks (weeks 1 to 8), the estimates jump sharply, demonstrating that the increase in hate crimes emerges only after and not before the pleasant-day protest.

Because the specification controls for municipality and week fixed effects as well as separate dynamic effects for pleasant weather and protest alone, this event-study pattern assuages several identification concerns. First, the flat pre-treatment path rules out differential pre-existing trends in hate crime across municipalities. Second, the lack of any anticipatory dips or spikes before week 0 suggests that neither would-be perpetrators nor protest organizers strategically time their actions around weather forecasts or expected violence. Third, the persistence of elevated estimates in the weeks following the protest underscores that our reduced-form captures genuine post-event emboldening.

Additional controls. Table B.2 reports how our preferred reduced-form estimate responds to the inclusion of an extensive set of additional controls, both in levels (Panel A) and interacted with the protest and weather dummies (Panel B). We report the precise sources and frequency of each of these variables in Appendix Table C.1. In Panel A, columns 1 through 5 sequentially add municipal covariates: GDP per capita, population density and unemployment; the share of refugees; right-wing AfD vote share; and overall crime rate. Across all specifications, the coefficient on the interaction of protest and pleasant weather remains tightly clustered around 0.086 with unchanged precision. Panel B repeats the exercise but interacts each new control with both protest occurrence and weather. After allowing for heterogeneous effects of these covariates, the coefficient on the interaction term becomes slightly larger in magnitude and hovers around 0.110 compared to 0.086 in the baseline specification. Together, these exercises show that neither baseline differences in municipal characteristics nor their interaction with weather or protest timing can explain away the marked increase in hate-crime incidence following high-salience PEGIDA demonstrations.

Augmenting our preferred specification with a rich set of municipal covariates and their interactions with protest occurrence and weather strengthens our identification by directly addressing omitted variable bias. If factors such as local economic conditions, political leanings, crime trends, refugee shares or even

concurrent protest activity were driving both turnout and hate crimes, then including them as controls would materially alter the estimated effect of pleasant-day protests on violence. The fact that our key coefficient remains essentially unchanged even once we allow these covariates to have heterogeneous effects depending on whether a protest took place or the weather was pleasant shows that none of these potential confounders can account for our finding. Moreover, interacting controls with both the protest dummy and the weather dummy captures any systematic way in which these municipal characteristics might amplify or dampen the impact of protest and weather alone on hate crimes. By doing so we guard against the concern that, say, wealthier areas react differently to protest or that regions with higher refugee shares are uniquely sensitive to weather conditions.

Reporting Bias. Table B.3 presents a series of placebo and reporting-bias checks. In Germany, policing budgets, resource allocations, and many operational protocols are set at the state level rather than by individual municipalities. State interior ministries decide how many officers to deploy, what training they receive, and how rigorously different offenses are investigated. By adding state-week (and region-week) fixed effects, we sweep out any contemporaneous shifts in state-level policing intensity or reporting standards whether driven by budget decisions, political directives, or public attention that coincide with high-salience protests. Therefore, in column 1 we add state by week fixed effects (there are 16 states in Germany) and in column 2 we replace these with region by week fixed effects (there are 38 NUTS-2 regions in Germany); in both cases the coefficient on the interaction term remains essentially unchanged at about 0.09, mitigating the concern that either state-level or regional shifts in policing or record-keeping can account for our main result.

Columns 3 through 5 then test for changes in crime reporting at the municipal level. Our hate crime data contains information on whether a suspect has been identified in the case, and distinguishes between three broad crime categories: rallies, arson and assault. We have shown in Figure 3 that rallies and other protest-related incidents cannot account for the surge in hate crimes. Here, we leverage the fact that some types of hate crimes are easy to observe and can be classified as hate crimes with little reporting or policing effort. Column 3 uses as its dependent variable an indicator for whether any recorded hate crime that week had an identified suspect; the coefficient shows no uptick in suspect identification after high-salience protests. Column 4 examines the ratio of hate crimes with suspects to total hate crimes has increased. Column 5 takes the share of arson incidents (arguably the most conspicuous offenses) as the outcome. Throughout, coefficients are small and statistically insignificant. Together, these findings rule out a story in which the surge in recorded hate crimes simply reflects heightened vigilance or reporting by police and media, reinforcing the interpretation that we observe a genuine increase in anti-minority violence.

Spillovers and spatial correlation. Our analysis is at the municipality level, which is the finest geographical level at which data is available to us. There may be spatial spillovers 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 3, we replicate our main analysis at higher levels of aggregation to account for those spillovers in Table B.4. We report our baseline estimate in column 1. In columns 2 and 3 of Table B.4, we collapse our dataset to higher geographical administrative levels (NUTS-3 and NUTS-2 instead of municipalities for column 2 and 3 respectively).³¹ NUTS-3 regions correspond to cities and their suburbs in most cases. NUTS-2 regions are quite broad and capture entire states in some cases. Our result remain precisely estimated and increase in magnitude.

In our main analysis standard errors are clustered at the municipality level: this implies that the errors of observations in the same municipality may be correlated, but not the errors of observations

³¹We collapse the data by taking the mean of the hate crime, pleasant, and protest indicators over municipalities. The interaction of protest and weather is then the interaction of the means.

from different municipalities. However, observations from neighboring municipalities may be correlated, for example if municipalities are exposed to similar events. While we show that geographic spill-overs are small in Table 3, we address this concern here by allowing for spatially correlated standard errors. In columns 4 to 7 of Table B.4, we use a panel version of Conley standard errors (Conley, 1999) allowing correlation between observations from municipalities at a given distance (50, 100, 150 and 200 km), even at different time periods. In all cases, the effect remains precisely estimated.

Alternative clustering structures. Table B.5 assesses whether our inference is sensitive to alternative assumptions about the correlation structure of the error term. Column 1 reproduces our baseline reduced-form estimate under one-way clustering at the municipality level. Columns 2 to 5 implement a set of increasingly conservative two-way clustering schemes that allow for common shocks across municipalities within the same calendar time unit, while still accounting for serial correlation within municipalities. Specifically, we double cluster by municipality and week (column 2), municipality and month (column 3), municipality and quarter (column 4), and, more coarsely, by district and month (column 5), thereby permitting spatially proximate municipalities within the same administrative district to share correlated shocks over a given month. Across these clustering choices, the standard errors increase only modestly. Columns 6 and 7 further relax the dependence structure by applying Driscoll-Kraay standard errors, which are robust to very general forms of cross-sectional dependence and serial correlation, using bandwidths of 4 and 13 weeks, respectively. Even under these assumptions, the effect remains precisely estimated, indicating that our baseline conclusions are not driven by a particular clustering choice and are robust to allowing for both spatial and temporal correlation in unobservables.

Variation on the fixed effect structure. Our main analysis includes municipality and week fixed effects. We show in Table B.6 how our core reduced-form estimate varies when we flexibly absorb different slices of unobserved heterogeneity that might otherwise confound the Protest \times Weather effect. Column 1 reproduces the baseline specification with week and municipality fixed effects plus our standard controls. In column 2 we replace the simple municipality dummies with municipality by month-of-year fixed effects. These absorb any annual seasonal patterns unique to each municipality, including summer festivals or seasonal economic activities, that could correlate both with pleasant-day turnout and with hate-crime reporting. Column 3 returns to the standard week and municipality fixed effects but adds a municipality-specific linear time trend. This captures any gradual, linear evolution in anti-minority violence or right-wing mobilization at the local level, perhaps reflecting slow shifts in demographics or long-run political realignment. Column 4 combines the month-of-year FE from column 2 with the municipality time trends of column 3, thereby sweeping out both non-linear seasonal patterns and linear secular trends. Finally, column 5 replaces week fixed effects with state by week fixed effects. By doing so, we allow for every state to experience its own week-to-week shocks such as state-level election campaigns, policing directives, or budgetary cycles that might influence both protest visibility and hate-crime incidence. Across all five specifications, the coefficient on the interaction of protest and pleasant weather remains large, positive, and precisely estimated. The stability of these estimates demonstrates that neither municipality-specific seasonality, local secular trends, nor state-level time shocks can explain away our main finding: protests that draw higher turnout under favorable weather conditions causally raise the likelihood of anti-minority violence.

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 (2023), 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 robustness of our results to these issue, we leverage the estimator from De Chaisemartin & D’Haultfoeuille (2022) to estimate the DiD_1 effect, i.e. the effect immediately after treatment by a protest that experienced pleasant weather. We present the results in Table B.7. Our units may be treated multiple times. Column 2 assumes that, in the period immediately following treatment, the effect of the pleasant protest becomes time-invariant (i.e. units treated at t have, starting at $t + 2$, parallel trends to untreated units). This is equivalent to treating treated units as “new” units for the purpose of fixed effects. Column 3 relaxes this assumption, and instead excludes observations for 5 weeks after a pleasant protest: the assumption is that, if no further pleasant protest happened, the trend is parallel to untreated units starting at $t + 5$ – this matches the graph of the event study in Figure B.1. Column 1 presents the baseline two-way fixed effects estimate for comparison. In both cases, we observe similar estimates that are precisely estimated, and slightly larger in magnitude from the two-way fixed effects. The change in magnitude is unsurprising, as alternative DiD estimators by design apply a different weighting to the per-unit treatment effects.

Nonlinearity, Poisson and logit. Our main specification uses a linear probability model to estimate the effect of protest on hate crimes. We check that our results are robust to using non-linear models. Results are presented in columns 4 and 5 of Table B.7. Column 4 estimates a logit model. In column 5, we perform a Poisson regression using the number of hate crimes as an outcome. In both cases, we obtain precisely estimated positive effects on, respectively, the probability of a hate crime occurring and the number of hate crimes. Table B.8 shows alternative definitions of the salience variables. Our baseline specifications uses the one plus count of turnout, newspaper reports and PEGIDA tweets to account for their skewed distribution. Each panel of Table B.8 transforms the outcome variable as the per capita count in column 2, the inverse hyperbolic sine transformed count in column 3, and the standardized count in column 4. With the exception of the per capita turnout and standardized newspaper articles which become less precisely estimated, all other estimates remain consistent in sign and precision.

Definition of the weather cut-offs. Our analysis relies on specific weather cut-offs. In Figure B.3, we examine the dependency of our results on the weather cut-off by plotting the estimates for different cutoffs. We decrease the precipitation threshold (mean rain during protest hours) by 1 to 7 mm. We decrease the maximum temperature threshold (maximum temperature during protest hours) by 1, 2 and 3 degrees, and do the same with the minimum temperature threshold. Finally, we increase the minimum temperature threshold then the maximum temperature threshold by 1, 2, and 3 degrees. The coefficients are statistically indistinguishable in size and similar in precision. Note also from Figure A.1 (which reports Kleinbergen-Paap F-statistics from our preferred specification for different pleasant weather definitions) that the predictive power for turnout drops significantly for minimum temperatures under -2 degrees Celsius, as well as maximum temperatures below 18 and above 22 degrees Celsius - which may explain why reduced form coefficients for these cut-offs in Figure B.3 are slightly smaller and less precisely estimated.

Robustness to outliers. In Figure B.4, we examine the possibility that our results are driven by some municipalities that react very strongly to large protests or by the fact 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 2SLS coefficients and confidence intervals of our treatment variable dropping one NUTS-3 region at a time. Our results are very similar in size and precision throughout, suggesting that no single NUTS-3 is driving the effect. The bottom panel shows the same estimates dropping single weeks, further confirming the robustness of our estimates.

Table B.1. **Robustness: Serial Correlation, Past Protests, and Seasonality**

| | Any hate crime [0,1] | | |
|---------------------------|----------------------|---------------------|-------------------------|
| | (1) | (2) | (3) |
| Protest \times Weather | 0.080*** (0.026) | 0.086*** (0.026) | 0.097*** (0.029) |
| Observations | 2,820,944 | 2,820,944 | 2,820,944 |
| Municipalities | 10825 | 10825 | 10825 |
| Week FE | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes |
| Protest \times Controls | Yes | Yes | Yes |
| Weather \times Controls | Yes | Yes | Yes |
| Additional controls | cumul. past protest | cumul. past turnout | muni \times m.o.y. FE |
| Mean dep. var. | 0.00330 | 0.00330 | 0.00330 |

Notes: This table presents results from the main estimating equation 1. Outcome is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Column 1 controls for the cumulative number of past PEGIDA Monday protests in the municipality up to week t . Column 2 additionally controls for the cumulative number of past PEGIDA Monday protest participants. Column 3 includes municipality-by-month-of-year fixed effects. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.2. **Additional Controls**

| | Any hate crime [0,1] | | | | |
|--|----------------------|---------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) |
| Panel A: Simple Controls | | | | | |
| Protest × Weather | 0.086*** (0.026) | 0.087*** (0.026) | 0.087*** (0.026) | 0.085*** (0.029) | 0.085*** (0.029) |
| Protest AME | 0.048** (0.021) | 0.048** (0.021) | 0.049** (0.021) | 0.050** (0.022) | 0.050** (0.022) |
| Panel B: Controls Interacted With Weather and Protest | | | | | |
| Protest × Weather | 0.086*** (0.026) | 0.109*** (0.026) | 0.111*** (0.026) | 0.106*** (0.029) | 0.105*** (0.029) |
| Protest AME | 0.048** (0.021) | 0.085** (0.039) | 0.082** (0.036) | 0.104*** (0.035) | 0.063 (0.043) |
| Observations | 2,820,944 | 2,809,879 | 2,781,952 | 2,781,843 | 2,781,843 |
| Municipalities | 10825 | 10784 | 10677 | 10677 | 10677 |
| Week FE | Yes | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes | Yes |
| Protest × Controls | Yes | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes | Yes |
| Economic | | Yes | Yes | Yes | Yes |
| Refugee | | | Yes | Yes | Yes |
| Right-wing | | | | Yes | Yes |
| Crime | | | | | Yes |
| Mean dep. var. | 0.00286 | 0.00286 | 0.00284 | 0.00281 | 0.00281 |

Notes: This table presents results from the main estimating equation 1 and successively adding controls in Panel A and their respective interactions with weather and protest in Panel B. The outcome is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Column 1 shows the baseline specification. Column 2 includes GDP per capita, population density and unemployment rate. Column 3 additionally adds the share of refugees in the municipality. Column 4 includes the vote for AfD as a measure of right-wing politics. Column 5 further controls for crime rate (per capita). All variables vary by municipality and year. At the bottom of the table we indicate the average marginal effect (AME) of protest when all controls are taken at their mean value, and the corresponding standard error in parentheses. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3. **Reporting Bias**

| | Any hate crime [0,1] | | any suspect | suspect ratio | arson ratio |
|--------------------|----------------------|---------------------|-------------------|----------------------|----------------------|
| | (1) | (2) | (3) | (4) | (5) |
| Protest × Weather | 0.093*** (0.030) | 0.092*** (0.030) | -0.012 (0.009) | -0.005 (0.005) | 0.003 (0.002) |
| Observations | 2,820,471 | 2,820,471 | 2,820,944 | 2,820,944 | 2,820,944 |
| Municipalities | 10823 | 10823 | 10825 | 10825 | 10825 |
| Municipality FE | Yes | Yes | Yes | Yes | Yes |
| Week FE | | | Yes | Yes | Yes |
| State-week FE | Yes | | | | |
| NUTS2-week FE | | Yes | | | |
| Protest × Controls | Yes | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes | Yes |
| Protest AME | 0.047** (0.022) | 0.047** (0.022) | 0.002 (0.005) | -0.002*** (0.000) | -0.002*** (0.001) |
| Mean dep. var. | 0.00276 | 0.00276 | 9.32e-05 | 7.98e-05 | 7.27e-05 |

Notes: This table presents results from the main estimating equation 1 and investigates bias in the reporting of hate crimes. The outcome in columns 1 and 2 is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Column 1 adds state-week fixed effects, and column 2 adds NUTS-2-week fixed effects. Column 3 reports hate crimes with at least one identified suspect the ratio. Column 4 reports the share of hate crimes with suspects to total hate crimes. Column 5 reports the share of arson attacks to total hate crimes. Hate crime data come from ProAsyl and AAF. At the bottom of the table we indicate the average marginal effect (AME) of protest when all controls are taken at their mean value, and the corresponding standard error in parentheses. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.4. **Robustness to Alternative Spatial Structure**

| | Any hate crime | | | | | | |
|--------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| | Analysis level | | | Spatial clustering | | | |
| | municipality (1) | NUTS-3 (2) | NUTS-2 (3) | 50 km (4) | 100 km (5) | 150 km (6) | 200 km (7) |
| Protest × Weather | 0.086*** (0.026) | 0.130*** (0.029) | 0.311*** (0.017) | 0.086*** (0.023) | 0.086*** (0.014) | 0.086*** (0.017) | 0.086*** (0.011) |
| Observations | 2,820,944 | 104,452 | 9,869 | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 |
| Geographic units | 10,825 | 401 | 38 | 10,825 | 10,825 | 10,825 | 10,825 |
| Week FE | Yes |
| Geographic unit FE | Yes |
| Protest × Controls | Yes |
| Weather × Controls | Yes |
| Mean dep. var. | 0.00330 | 0.0280 | 0.0361 | 0.00330 | 0.00330 | 0.00330 | 0.00330 |

Notes: This table presents results from the main estimating equation 1 but varying the geographic units in columns 1 to 3 (municipality, NUTS-3 equivalent to districts, NUTS-2 equivalent to commuting zones) and accounting for spatially correlated standard errors in columns 4 to 7 (between 50 and 200 kilometers). There are approximately 10,000 municipalities, 401 NUTS-3 regions, and 38 NUTS-2 regions. Outcome is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Column 1 controls for the cumulative number of past PEGIDA Monday protests in the municipality up to week t . Column 2 additionally controls for the cumulative number of past PEGIDA Monday protest participants. Column 3 includes municipality-by-month-of-year fixed effects. Standard errors are clustered at the municipality level. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.5. **Alternative Clustering Structures**

| | Any hate crime [0,1] | | | | | | |
|--------------------|----------------------|---------------------|---------------------|----------------------|-----------------------|--------------------------|-------------------|
| | baseline (1) | double clustering | | | | Driscoll-Kraay std. err. | |
| | | muni. week (2) | muni. month (3) | muni. quarter (4) | district month (5) | 4 weeks (6) | 13 weeks (7) |
| Protest × Weather | 0.086*** (0.026) | 0.086*** (0.028) | 0.086*** (0.029) | 0.086*** (0.041) | 0.086* (0.029) | 0.086* (0.045) | 0.086* (0.044) |
| Observations | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 |
| Municipalities | 10825 | 10825 | 10825 | 10825 | 10825 | 10825 | 10825 |
| Week FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Protest × Controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.00286 | 0.00286 | 0.00286 | 0.00286 | 0.00286 | 0.00286 | 0.00286 |

Notes: This table presents results from the main estimating equation 1 but varying the clustering methods. Outcome is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Column 1 uses our baseline clustering strategy by municipality. Column 2 uses two-way clustering by municipality and week, column 3 by municipality and month, column 4 by municipality and quarter, column 5 by district and month. Column 6 use Driscoll-Kraay standard errors with a Bartlett kernel with a 4 weeks bandwidth, and column 7 with a 13 weeks bandwidth. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6. **Alternative Fixed Effects Specifications**

| | Any hate crime [0,1] | | | | |
|------------------------------|----------------------|---------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) |
| Protest \times Weather | 0.086*** (0.026) | 0.097*** (0.029) | 0.088*** (0.027) | 0.100*** (0.033) | 0.093*** (0.030) |
| Observations | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,944 | 2,820,471 |
| Municipalities | 10825 | 10825 | 10825 | 10825 | 10823 |
| Week FE | Yes | Yes | Yes | Yes | |
| Municipality FE | Yes | | Yes | | Yes |
| Protest \times Controls | Yes | Yes | Yes | Yes | Yes |
| Weather \times Controls | Yes | Yes | Yes | Yes | Yes |
| Municipality \times MoY FE | | Yes | | Yes | |
| Week \times State FE | | | | | Yes |
| Municipality TT | | | Yes | Yes | |
| Protest AME | 0.048** (0.021) | 0.039* (0.020) | 0.041** (0.020) | 0.034* (0.021) | 0.047** (0.022) |
| Mean dep. var. | 0.00286 | 0.00286 | 0.00286 | 0.00286 | 0.00276 |

Notes: This table presents results from the main estimating equation 1 but varying the fixed effects structure. Outcome is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Column 1 reports the baseline specification. Column 2 replaces municipality FE with municipality-by-month-of-year fixed effects. Column 3 includes municipality-specific linear time trends. Column 4 replicates column 2 and adds municipality time trends. Column 5 replaces week fixed effects with state-by-week fixed effects. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.7. **Robustness to Alternative Estimation Methods**

| | Any hate crime [0,1] | | | | Count |
|---------------------------|----------------------|---------------------|-------------------|-------------------|---------------------|
| | baseline (1) | C&H t+2 (2) | C&H t+5 (3) | logit (4) | Poisson (5) |
| Protest \times Weather | 0.086*** (0.028) | 0.108*** (0.032) | 0.066* (0.040) | 0.788* (0.407) | 0.577*** (0.209) |
| Observations | 2,820,944 | 2,820,944 | 2,819,951 | 2,820,944 | 2,820,944 |
| Municipalities | 10825 | 10825 | 10825 | 10825 | 10825 |
| Week FE | Yes | Yes | Yes | Yes | Yes |
| Municipality FE | Yes | Yes | Yes | Yes | Yes |
| Protest \times Controls | Yes | Yes | Yes | Yes | Yes |
| Weather \times Controls | Yes | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.00286 | 0.00286 | 0.00286 | 0.00286 | 0.00339 |

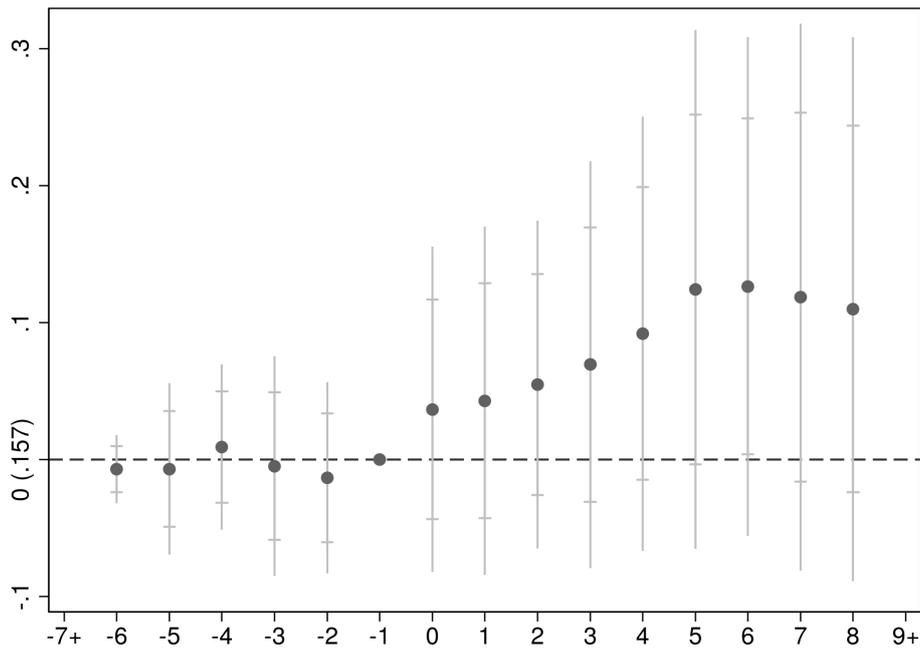
Notes: This table presents results from the main estimating equation 1 but varying the estimation method. Outcome is the probability of a hate crime against refugees in the six days following the protest i.e., between Tuesday and Sunday of the same week. Column 1 presents the baseline two-way fixed effects estimate. Column 2 uses the DiD₁ estimator from De Chaisemartin & D’Haultfoeuille (2022), assuming that the effect of pleasant protest becomes constant starting from the second period after treatment. Column 3 instead assumes that the effect of a pleasant protest becomes constant after 5 weeks. Column 4 uses a logit model. Column 5 uses a Poisson model with the number of hate crimes as outcome. The sample covers January 2015 to December 2019 at the municipality-week level. Data on protests and participants are from Kanol & Knoesel (2021); hate crimes originate from Pro Asyl and the Amadeu Antonio Foundation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.8. **Robustness to Alternative Outcome Definitions**

| | log(1+count) (1) | per 1000 (2) | IHS (3) | standardized (4) |
|---|---------------------|--------------------|---------------------|---------------------|
| Panel A: Protest turnout | | | | |
| Protest × Weather | 0.388*** (0.081) | 1.326 (1.197) | 0.391*** (0.082) | 0.353*** (0.075) |
| Observations | 2,820,944 | 2,809,931 | 2,820,944 | 2,820,944 |
| Municipalities | 10,825 | 10,784 | 10,825 | 10,825 |
| Panel B: Newspaper reports on PEGIDA protest | | | | |
| Protest × Weather | 0.337*** (0.096) | 0.031** (0.015) | 0.403*** (0.112) | 0.301 (0.209) |
| Observations | 2,820,944 | 2,809,931 | 2,820,944 | 2,820,944 |
| Municipalities | 10,825 | 10,784 | 10,825 | 10,784 |
| Panel C: PEGIDA tweets | | | | |
| Protest × Weather | 0.231** (0.116) | 0.052** (0.021) | 0.260* (0.133) | 0.158** (0.067) |
| Observations | 2,762,218 | 2,751,205 | 2,762,218 | 2,762,218 |
| Municipalities | 10,599 | 10,558 | 10,599 | 10,558 |
| Municipality & Week FE | Yes | Yes | Yes | Yes |
| Protest × Controls | Yes | Yes | Yes | Yes |
| Weather × Controls | Yes | Yes | Yes | Yes |

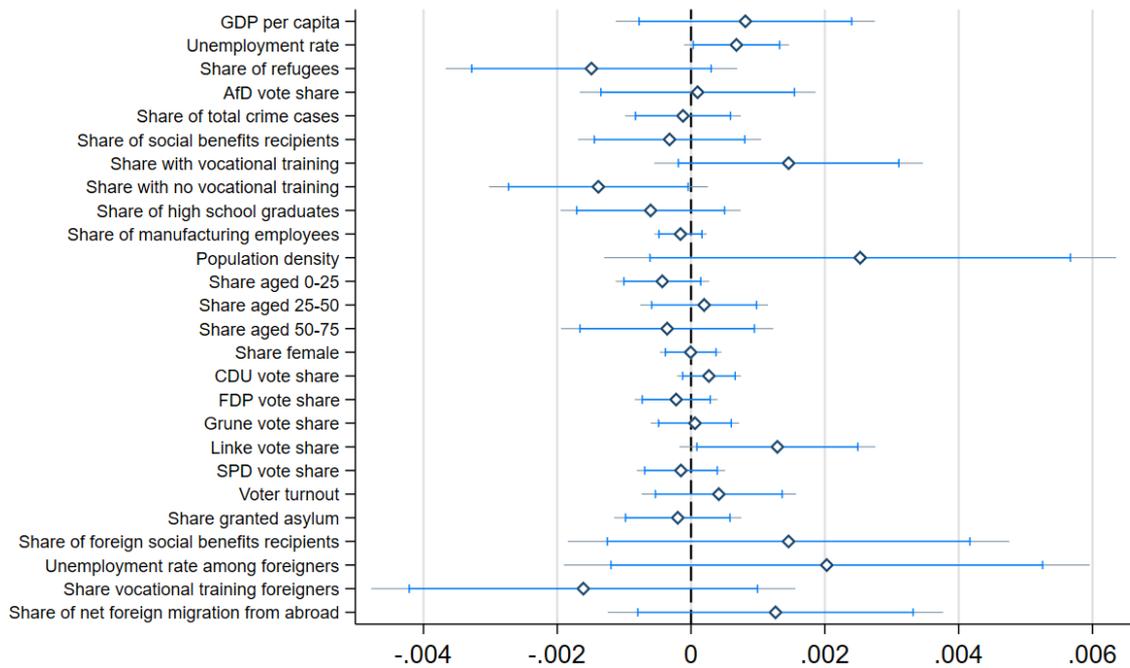
Notes: This table examines the robustness of the relationship between pleasant weather and protest salience to alternative definitions of the outcome variables. Each cell reports the coefficient on the interaction between a scheduled PEGIDA protest and pleasant weather. Panel A uses protest turnout as the dependent variable, Panel B uses the number of newspaper articles covering the protest, and Panel C uses the number of PEGIDA-related tweets. Columns differ only in the transformation of the dependent variable. Column (1) reports the baseline specification using log(1 + count), which is the specification used in the main text. Column (2) measures outcomes per 1,000 residents. Column (3) applies the inverse hyperbolic sine (IHS) transformation. Column (4) uses standardized outcomes constructed using the mean and standard deviation among municipality-weeks with protests. The sample covers January 2015 to December 2019 at the municipality-week level. All regressions include municipality and week fixed effects as well as interactions of protest and weather indicators with the full set of controls used in the baseline specification. Data on protests and participants are from Kanol & Knoesel (2021). Newspaper coverage is measured using the universe of newspaper articles from GENIOS and PEGIDA tweets are constructed from the universe of geo-localized tweets containing the term PEGIDA. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.1. Event Study: Hate Crimes Following Protests on Pleasant Days



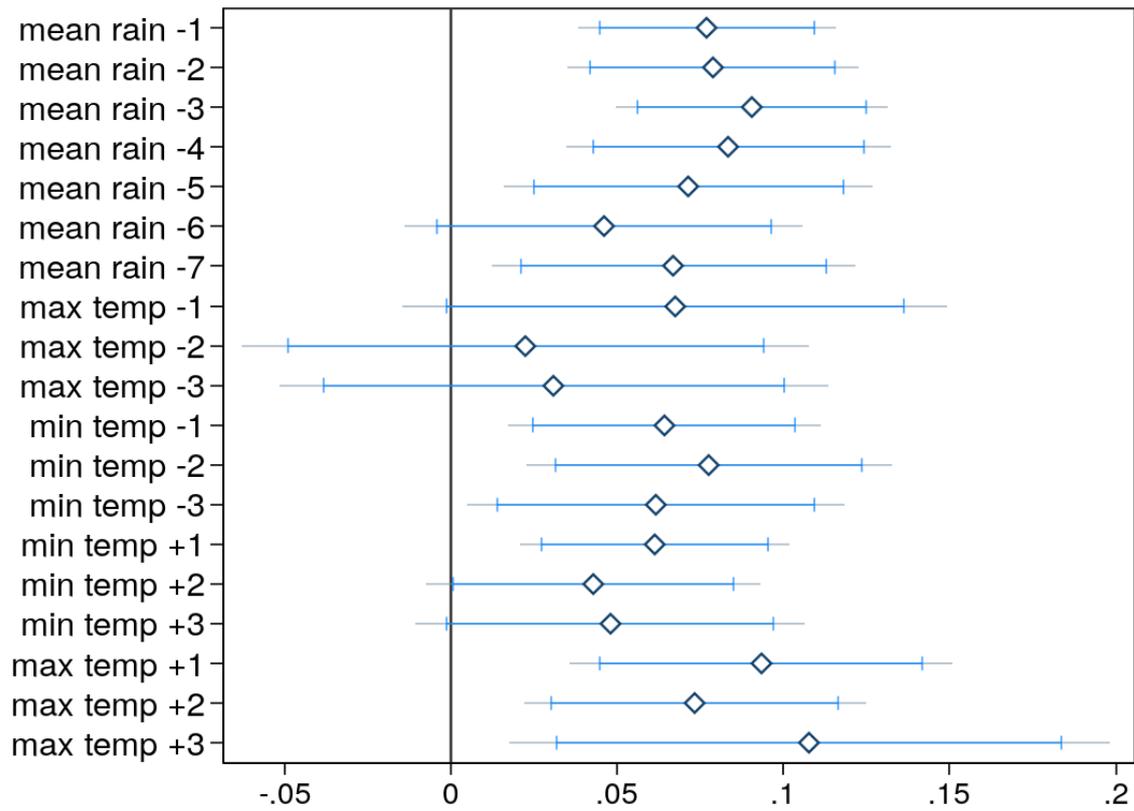
Notes: Event study using `xtevent` from Freyaldenhoven et al. (2024). Regression based on estimating equation $H_{it} = \alpha_i + \lambda_t + \sum_{k=T_0}^{T_1} \beta_k (\text{weather}_{ik} \times \text{protest}_{ik}) + \gamma_k(\text{weather}_{ik}) + \phi_k(\text{protest}_{ik}) + \epsilon_{it}$ at the municipality i and week level t where the outcome is a dummy for any hate crime. We plot β_k coefficients with 90 and 95 percent confidence intervals. Standard errors are clustered at the municipality level. The sample consists of ever-treated municipalities only.

Figure B.2. Balance: Municipality characteristics and protest on pleasant days



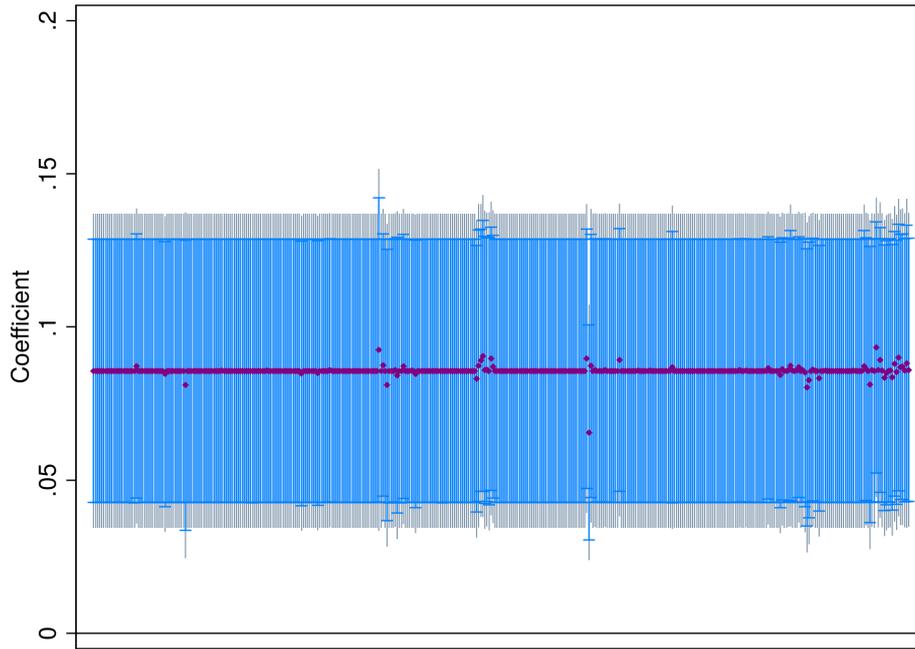
Notes: Number of pleasant-day protests in a given municipality and year regressed on various economic, political and socio-demographic municipality characteristics between 2015 and 2019 from the INKAR data base. Regression includes state and year fixed effects, standard errors are clustered on the state level, controls include the total number of pleasant days and the total number of protests in a given municipality and year.

Figure B.3. Effect of Weather \times Protest using alternative cut-offs for pleasant weather dummy

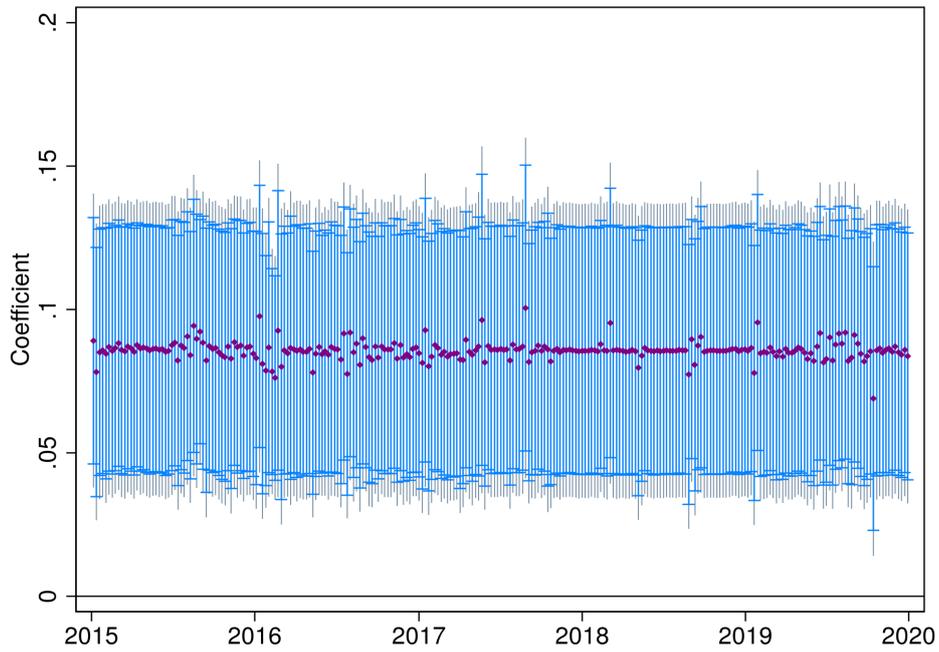


Notes: Coefficient plot of pleasant-day Monday PEGIDA protest with 90 and 95 percent intervals of the main estimating equation 1, varying the definition of pleasant weather. Each line varies one characteristic of the pleasant weather indicator while keeping the other constant. Lines 1-7 decrease the threshold of rain by 1 to 7 mm respectively. Lines 8-10 (17-19) vary the max temperature by (-)1, (-)2 and (-)3 degrees Celsius respectively. Lines 11-16 vary the min temperature threshold.

Figure B.4. Robustness to excluding potential outliers



(a) Dropping NUTS-3 regions



(b) Dropping weeks

Notes: Coefficient plot of pleasant-day Monday PEGIDA protest with 90 and 95 percent intervals of the main estimating equation 1, varying the definition of pleasant weather, dropping one NUTS-3 region at a time in the top panel; and dropping one week at a time in the bottom panel. There are a total of 401 NUTS-3 regions and 260 weeks.

Appendix C: Data Appendix

This Appendix describes the underlying data in more detail. We summarize the main variables, their sources, time span and geographic coverage in Table C.1.

C.1 PEGIDA Protest Data

To create this dataset, 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 $i, j \in [\text{Göttingen, Osterode am Harz}]$ 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 Terrorist attacks.

To study the extensive margin of protest, we use Islamist terrorist attacks as an exogeneous source of variation. We use the Global Terrorism Database (START, 2022), and extract all terrorist attacks that occurred in the European Union, the United Kingdom, Norway or Switzerland, excluding those attacks that occurred in Germany. We select only attacks that are attributed to Islamist groups or individuals, and construct a variable equal to 1 if there was such an attack in the last 28 days.

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.³² 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 2021). 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 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,³³ 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

³²Since 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/>.

³³Due to the restriction on the Twitter API placed in March 2023, we are unable to complete this dataset to the end of 2019.

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.

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 by selecting all articles published between 2015 and 2019 that contain a word ending in "GIDA" (e.g. PEGIDA, THÜGIDA, ...) or a word in relation to immigrants or refugees and a word related to demonstrations. We then use a large language model (gemini-2.0-flash-001) with a custom prompt to determine whether the articles mention pro-PEGIDA protests or counter-protests. We require that the output is given in JSON format following a given schema, and set the generation temperature to 0.³⁴ 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 replaced 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
(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"]
}
```

³⁴The temperature introduces randomness in the LLM output. Setting the temperature to zero ensures reproducibility.

```
}  
}
```

We then re-process the results to obtain a protest dataset. First, we manually reclassify the demonstration type, as the model doesn't always simply respond "PEGIDA" or "anti-PEGIDA", but sometimes uses other denominations ("THÜGIDA"), merges demonstrations ("both"), or includes unrelated demonstrations that are mentioned in the article ("pro-Hong Kong"). When "both" is indicated, we expand the observation into two protests with missing number of participants.

Second, we process the location. We first remove locations that are clearly ambiguous: sometimes, the model doesn't respond with a city, but a location in the city (e.g. "Altmarkt" for "Altmarkt, Dresden"). We also remove locations that are too broad ("Bavaria", "all over Germany"). Filtering these out is important, as they could match other locations in Germany. We then use Nominatim to geocode the locations to municipalities, and remove locations outside Germany (e.g. some demonstrations took place in Austria).

In the date field, the model will sometimes instead answer with imprecise dates ("early 2016"), sometimes matching the requested format ("2016/01/00"). We filter out these date. We also remove dates in the future (sometimes articles announce the expected number of participants), and dates too far into the past (more than 60 days) as they correspond to often less precise background information given in the articles.

Finally, we process the number of participants. Here again, the model sometimes answers with vague textual descriptions matching the article, such as "more than a thousand". In this case, we manually process the result and use the lower end of the provided range.

In order to transform the dataset of coverage of protest into a dataset of protest, we collapse observations of the same type of demonstration (pro or anti-PEGIDA) in the same city at the same date, and take the median of the reported numbers of 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 newspapers (or more precisely advertisement placement units³⁵) 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.³⁶ We end up matching 485 advertising units, representing 84% of the total measured diffusion, to 134 publications referenced 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 model³⁷, we train the LLM to carefully

³⁵The 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.

³⁶The 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.

³⁷The 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.

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",
            "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 if mentioned
- (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.³⁸

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.

³⁸English 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.

Not all hate crimes in the dataset have a complete description: in particular, crimes reported in response to parliamentary inquiries are not always matched to news reports containing more information, and instead contain a generic description. Out of the 11 618 hate crimes in the dataset, we have 2841 unique descriptions, representing about 25% of the entries.

Table C.1. Description of Data Sources

| Variable | Regional Level | Period | Source |
|---|----------------|--------------|--|
| Main variables | | | |
| PEGIDA protests and number of participants | 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) |
| Media variables | | | |
| Sample of German language tweets | dist | 2013-2018 | Twitter API |
| Tweets mentioning PEGIDA | dist | 2015-2020 | Twitter API |
| #RefugeesWelcome tweets | 2014-2018 | baseline | Twitter API |
| Counter protest | muni | 2015-2020 | Processing of GENIOS articles |
| Articles about pro and anti-PEGIDA protests | muni | 2015-2020 | Processing of GENIOS articles |
| Diffusion of newspapers | muni | 2020 | Informationsgemeinschaft zur Feststellung der Verbreitung von Werbeträgern |
| Additional 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 |
| Asylum recipients (share) | dist | 2011-2021* | Federal Statistics Office |
| AfD votes (Bundestag) | muni | 2013-2021*** | Federal Returning Officer |
| Total crime cases (per 100k pop) | muni | 2013-2021 | Federal Criminal Police Office |

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. Column 3 provides the source from which we extract each of the variables.

Table C.2. Examples of Hate Crimes

| Hate Crime | Description |
|------------|--|
| Assault | <i>At a gas station on Linxweilerstraße, a man from Syria was attacked and injured around 9:15 PM, allegedly for racist reasons. According to the police, a young man first insulted the victim by calling him a "Kanacke" and then struck him on the nose with his forearm. The police describe the perpetrator as 20 years old, approx. 1.70 meters tall, of strong build, with short, medium-blond hair and a beard.</i> |
| Arson | <i>A paper container and doormats were set on fire at an asylum shelter, causing property damage. The Munich General Prosecutor's Office and the Lower Bavaria Police Headquarters issued a press release in early January, stating that a 42-year-old man from the Dingolfing-Landau district was arrested as a suspect in mid-December. He has been in custody since then. Preliminary analysis of the suspect's intercepted chat communications suggests that he set the fire with the intent to kill residents of the shelter out of xenophobic motivation.</i> |
| Rally | <i>Late in the evening, around 15 to 20 suspected neo-Nazis marched through the town with torches before lighting a larger fire in the church square. Their unannounced appearance, dressed in dark clothing, carrying torches, and wearing white masks, is reminiscent of previous actions by the self-proclaimed "Immortals." The local police also believe the action was politically motivated. According to the Mitteldeutsche Zeitung, some participants could be clearly identified as belonging to the far-right, and some were already known to the police for their propensity for violence.</i> |

Note: Examples of hate crimes classified as assault, arson, and intimidation.

Table C.3. Islamist Terrorist Attacks Used for the Extensive Margin Analysis

| Date | Country | Date | Country | Date | Country |
|------------|----------------|------------|----------------|------------|----------------|
| 2015-01-07 | France | 2016-09-01 | Denmark | 2017-08-18 | Spain |
| 2015-01-08 | France | 2016-09-04 | France | 2017-08-25 | Belgium |
| 2015-01-09 | France | 2016-09-10 | Spain | 2017-09-15 | France |
| 2015-02-14 | Denmark | 2016-10-05 | Belgium | 2017-09-15 | United Kingdom |
| 2015-04-19 | France | 2016-12-23 | Italy | 2017-09-30 | France |
| 2015-06-26 | France | 2017-02-03 | France | 2018-03-12 | Austria |
| 2015-08-21 | France | 2017-03-18 | France | 2018-03-23 | France |
| 2015-09-29 | Denmark | 2017-03-22 | United Kingdom | 2018-05-12 | France |
| 2015-11-13 | France | 2017-04-07 | Sweden | 2018-05-29 | Belgium |
| 2015-11-17 | United Kingdom | 2017-04-08 | Norway | 2018-08-20 | Spain |
| 2015-12-05 | United Kingdom | 2017-04-20 | France | 2018-08-31 | Netherlands |
| 2016-01-01 | France | 2017-05-18 | Italy | 2018-10-06 | Greece |
| 2016-01-07 | France | 2017-05-22 | United Kingdom | 2018-11-20 | Belgium |
| 2016-01-11 | France | 2017-06-03 | United Kingdom | 2018-12-11 | France |
| 2016-02-14 | Sweden | 2017-06-06 | France | 2018-12-31 | United Kingdom |
| 2016-02-18 | United Kingdom | 2017-06-07 | United Kingdom | 2019-03-05 | France |
| 2016-03-22 | Belgium | 2017-06-19 | France | 2019-03-18 | Netherlands |
| 2016-03-24 | United Kingdom | 2017-06-20 | Belgium | 2019-05-24 | France |
| 2016-05-26 | France | 2017-06-30 | Austria | 2019-08-31 | France |
| 2016-06-13 | France | 2017-08-05 | France | 2019-10-03 | France |
| 2016-07-14 | France | 2017-08-09 | France | 2019-11-29 | United Kingdom |
| 2016-07-20 | United Kingdom | 2017-08-17 | Spain | | |
| 2016-07-26 | France | 2017-08-18 | Finland | | |

Note: List of date and location of terrorist attacks used for the extensive margin analysis. Extracted from START (2022), including all attacks in the European Union, United Kingdom, Switzerland and Norway, but excluding Germany.