

Supplemental Data Appendix for Paper:
Terrorism and Voting:
The Rise of Right-Wing Populism in Germany

Navid Sabet* Marius Liebold Guido Friebe

This Version: January 27, 2025

Our paper aims at understanding the impact of terrorist attacks on right-wing voting in Germany. We compiled a new dataset from several different sources to carry out this study, which we describe more fully in this Appendix.

The Global Terror Database

We collect data on terror attacks in Germany between 2010 and 2020 from the Global Terror Database (GTD 2020) which is maintained by the University of Maryland, College Park. This is an open source database that documents information on terror attacks from around the world from 1970 to the present day. The database is maintained through data collection efforts from public, unclassified materials including media articles and electronic news archives, existing datasets and secondary source materials such as legal documents and books.

For an event to be included in the GTD several criteria must be met. First, the incident must be intentional, it must entail some level of violence and it must be perpetrated by sub-national actors. In other words, the database does not include state-sponsored acts of terrorism. Second, two of the following criteria must also be met: (i) The act must be aimed at attaining a political, economic, religious, or social goal; (ii) there must be evidence of an intention to coerce, intimidate, or convey some other message to a larger audience beyond the immediate victims; and/or (iii) the incident must occur outside the context of legitimate warfare.

The GTD provides information with regard to the identity of the target and the motivation of the perpetrator(s), though the latter information is not always complete. We therefore complete this information by looking up each of the 232 attacks using our news data and the internet to obtain information on the identity of the perpetrator and the motives behind the attack (Google 2024; LexisNexis 2025). Doing so enables us to classify 211 of the 232 attacks. The majority of the attacks (116 of the 211, or 55 percent) are carried out in the name of right-wing extremist causes and 57 percent target non-Germans, in line with the example illustrated in Section 2 of the main manuscript. If, however, we consider only the 124 first attacks in each of the unique 124 municipalities targeted by an attack, the figures are considerably higher: 75 percent of these attacks are carried out by right-wing extremists and 75 percent target foreigners.

*Corresponding author at: Faculty of Economics and Business, Goethe University Frankfurt, Theodor-W.-Adorno-Platz 4, 60629 Frankfurt, Germany. Email: sabet@econ.uni-frankfurt.de. Telephone: +49 (69) 798-34803.

NSADP Vote Share in 1933

Falter and Hänisch 1990 provide data on the *National Socialist German Workers' Party* (NS-DAP) vote share for federal elections between 1920 and 1933 in interwar Germany. This source includes the major parties' municipality-level election results for cities and villages with more than 2000 inhabitants during the 1920s. As no common identifier exists between this source and the data provided by *the Bundeswahlleiter*, we have to match municipalities between the sources.

We use `MLMATCH`, a novel state-of-the-art EM framework introduced by Karapanagiotis and Liebold 2023 to link municipalities across sources. This framework uses a similarity encoder to translate pairs of entity records from both data sources to machine learning-compatible numeric data. When translating the data, the encoder accounts for a range of similarity concepts (e.g., Levenshtein, Jaro-Winkler, Euclidean, etc.). Moreover, aside from using various similarity functions on one pair of variables, it also allows us to account for many variable combinations. The encoders' output is fed to two layers of deep artificial neural networks, resulting in a matching probability for each potential combination of records from both sources.

In our setup, we use nine similarity functions when comparing the non-harmonized municipality names between the two sources. Aside from this, we rely on the networks' standard specifications. As illustrated by Column 1 of Table 1, we train the model on 80% of the manually labeled subsample. The subsample corresponds to the municipalities in which terror attacks occurred during our observation period and for which we could find a match between the data sources.

Table 1
Benchmark Performance

(1)	(2)	(3)	(4)	(5)	(6)
Dataset	Fraction	Accuracy	Precision	Recall	F1-Score
Training	0.8	99.7	98.7	93.8	96.2
Testing	0.2	98.4	100	96.9	98.4

Note: This table illustrates `MLMatch`'s performance when matching historic municipalities from Falter and Hänisch 1990 with their current representations obtained from *the Bundeswahlleiter*. Column 1 indicates the dataset type underlying the performance evaluation. Column 2 indicates each dataset's fraction of the overall number of manually labeled municipalities. Columns 3 to 6 illustrate the corresponding performance metrics, including accuracy ($= \frac{TP+TN}{TP+TN+FP+FN}$), precision ($= \frac{TP}{TP+FP}$), recall ($= \frac{TP}{TP+FN}$), and the F-Score ($= 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$).

Columns 3 to 6 of Table 1 highlight the model's performance. When making predictions on the previously unseen testing data, the model does not lead to false positive matches, as indicated by a precision of 100. Moreover, the model likewise only makes very few false negative predictions, leading to a recall of 96.9 and a f1-score of 98.4.

News Coverage Data

We use news data from two sources. The first source, the *Frankfurter Allgemeine Zeitung* (FAZ), offers a country-wide and geographical unbiased news coverage. The FAZ is the second-largest national daily newspaper in Germany, measured in sales volumes. Unlike the tabloid press, the public considers the FAZ to sell high-quality journalism. The second data source is *LexisNexis* (LexisNexis 2025). LexisNexis provides access to news articles from

more than 1,000 newspapers through its online portal for academic research. Aside from a few more prominent news sources, LexisNexis predominantly covers local and regional newspapers, focusing on different geographic areas. The dataset construction is divided into three steps, outlined in the following.

Raw Data Collection & Transformation

For each terror attack in the GTD from our observation period in Germany, we collected news articles potentially covering the incident. We assume a news article to potentially cover a terror attack if it was published within the first ten days after the attack and if it included the name of the city of the attack. For *LexisNexis* data, we additionally established the criterion that a news article must incorporate at least one (case-insensitive) terror attack-related keyword (i.e., "Attacke", "Anschlag", or "Terror"). This additional condition reflects (opportunity) costs associated with obtaining *LexisNexis* data as the online portal only allows for semi-automated access. In total, we gathered approximately 105,000 news articles published by the *FAZ*, and 60,000 provided by *LexisNexis*.

Classification

Subsequently, to better identify news articles related to instances of terror, we implemented a three-step procedure. First, according to predictions by a classification model, we labeled the news articles as either related or unrelated to terrorism. In fact, we trained multiple classification models utilizing various classifier technologies and selected the best performing among them. Figure 2 summarizes the prediction quality of the models trained. As the MLP (Multilayer Perceptron) Classifier delivers the most robust results in terms of Accuracy (85.1 percent), we used this model for classifying the news articles. All models are trained on the same data, consisting manually labeled news articles collected from *LexisNexis* covering terror attacks in Austria and Switzerland between 2006 and 2018. The usage of Swiss and Austrian data is appropriate for two reasons. On the one hand, the language in Germany, Austria, and large parts of Switzerland, and thus of most news articles reporting on the incident, is the same. On the other hand, German newspapers frequently report on attacks in Germany's neighboring countries. We implemented the steps using the *scikit-learn* library for *Python*. Before we fit the model, we conducted preprocessing for both the training (372 news stories) and test data (94 reports). Specifically, we harmonized the spellings, performed lemmatization, and only kept lemmas representing adjectives, nouns, verbs, or institutions. During the preprocessing, we relied on the open-source software library *Spacy* and its associated *de_core_news_sm* pipeline, which comes pre-trained on German news data.

To minimize false-positive classifications, we establish a second condition for news articles to be considered terrorism-related. Articles must include at least one terrorism-related keyword. The list of valid keywords consists of the expressions "Terror", "Attacke", "Anschlag", "Bombe", "Messer", "Sprengstoff", "Blutbad", "Molotow", and "Attentat" in any style of writing. The third step guarantees that news articles labeled as related to terrorism indeed report on the terror attack of interest and not terrorism in general or a different incident. Accordingly, we manually check the classifications of all news articles previously labeled as terrorism-related. In summary, the implemented three-step procedure de facto eliminates false positive classifications and assignments. It, on the contrary, can not entirely rule out the existence of false-negative classifications. However, the confusion matrix of the trained MLP model (Figure 1) suggests that the problem of false-negative classifications is of minor relevance.

Table 2
Classification Model Performances

Model	Accuracy Score	Precision Score	Recall Score	F1 Score	Training Articles	Test Articles
SVC	0.606383	0.550000	0.536585	0.543210	372	94
KN	0.436170	0.436170	1.000000	0.607407	372	94
NB	0.765957	0.672727	0.902439	0.770833	372	94
DT	0.808511	0.810811	0.731707	0.769231	372	94
LR	0.808511	0.848485	0.682927	0.756757	372	94
RF	0.797872	0.958333	0.560976	0.707692	372	94
MLP	0.851064	0.813953	0.853659	0.833333	372	94

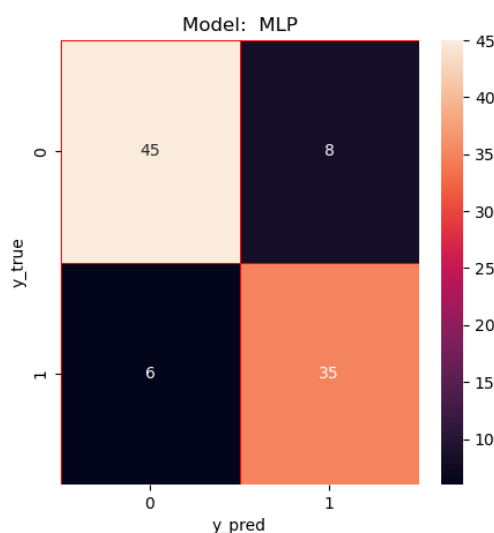


Figure 1
MLP Classification Confusion Matrix

Note: This figure plots the confusion matrix for the MLP classification model when evaluation the predictions of the test data.

Sentiment Analysis

When generating sentiment scores for news articles, we calculated the mean value of connotations associated with the words included in a report. The connotations capture the words' polarity (i.e., whether it has a positive or negative character), range between -1 (negative) and 1 (positive), and stem from the *SentiWS* word collection (Remus, Quasthoff, and Heyer 2010). For instance, whereas the term "hinterlistig" ("insidious") corresponds to a negative sentiment of -0.3187, "großzügig" ("generous") represents the positive score of 0.2077. The word collection contains sentiments of circa 3,500 base words and their 35,000 inflections.

SOEP Questions

We use data from the German Socio-Economic Panel (SOEP) which is a panel of individuals and households over time in order to study the political preferences and attitudes of the *same person* before and after an attack. We obtained access to the restricted-use SOEP data with municipality identifiers in order to link our data on successful/failed attacks to this survey data. Below, we provide details on the exact formulations of the survey questions used in the SOEP data and how we used them in our analysis.

Political attitude

We use variable *plh0004* to identify a person's position on the left-right political spectrum. Details are as follows:

Question:	In politics people sometimes talk of left and right. Where would you place yourself on a scale from 0 to 10 where 0 means the left and 10 means the right?
Answer Range:	An index from 0 (left) to 10 (right).
Usage:	We construct two binary variables using this measure. First, we generate a variable that is 1 if <i>plh0004</i> ≥ 6 and second, we construct an indicator that is 1 if <i>plh0004</i> ≥ 7 and 0 otherwise.

Party preferences

We use variable *plh0012_v6* to identify a person's political party preferences.

Question:	Which party do you lean toward?
Answer Range:	Categorical variable for all main parties in Germany.
Usage:	We construct multiple binary variables using this measure. First, we generate a variable that is 1 if <i>plh0012_v6</i> == <i>AfD</i> and 0 otherwise. Second, we construct an indicator that is 1 if <i>plh0012_v6</i> == <i>CDU</i> , <i>plh0012_v6</i> == <i>CSU</i> , or <i>plh0012_v6</i> == <i>CDU/CSU</i> and 0 otherwise. Subsequently, we repeat the exercise and label someone to prefer the social democrats if <i>plh0012_v6</i> == <i>SPD</i> , the left if <i>plh0012_v6</i> == <i>Die Linke</i> , the FDP if <i>plh0012_v6</i> == <i>FDP</i> , the green party if <i>plh0012_v6</i> == <i>Buendnis90/Gruene</i> . Moreover, we create a binary variable indicating ultra-rightwing party preferences that equals 1 if <i>plh0012_v6</i> == <i>NPD/Republikaner/DieRechte</i> and 0 otherwise.

Worried about migration to Germany

We use variable *plj0046* to identify people's concerns about migration to Germany.

Question:	How concerned are you about immigration to Germany.
Answer Range:	Range: 1 (very concerned), 2 (somewhat concerned), 3 (not concerned).
Usage:	We construct a binary variable that is 1 if a person is very concerned about migration to Germany and 0 otherwise.

Worried about (global) terrorism

We use variable *plh0039* to identify people's concerns about global terrorism.

Question:	How worried are you about global terrorism?
Answer Range:	1 (very concerned), 2 (somewhat concerned), 3 (not concerned at all).
Usage:	We construct a binary variable that is 1 if people are very concerned or somewhat concerned about global terrorism (i.e. $plh0039 == 1$ or $plh0039 == 2$) and 0 otherwise.

Political Participation

We use variable *pli0097_h* to identify people's participation in local political affairs.

Question:	Now some questions about your leisure time. Please indicate how often you take part in each activity: daily, at least once per week, at least once per month, seldom or never? (One of the listed activity reads out: Participating in political parties, municipal politics, citizens' initiatives)
Answer Range:	1 (daily), 2 (at least once a week), 3 (at least once a month), 4 (seldom), and 5 (never).
Usage:	We construct a binary variable that is 1 if people participate at least one time per month in local politics (i.e. $pli0097_h == 1$, $pli0097_h == 2$, or $pli0097_h == 3$) and 0 otherwise. Additionally, if a person residing in a municipality targeted by a terror attack has always participated in local political affairs according to the definition above before the attack, we label him/her politically active.

Other Variables

Variable	SOEP Label
Income	ijob1
Married	d11104
Female	d111021l
Uni Degree	bex8cert
Employed	e11102
Age	syear & birth_year
Moved	gkz
Rural area	kr_population & kr_area

References

- Falter, Jürgen W., and Dirk Hänisch. 1990. *Election and Social Data of the Districts and Municipalities of the German Empire from 1920 to 1933*. GESIS Data Archive, Cologne. ZA8013 Data file Version 1.0.0, <https://doi.org/10.4232/1.8013>. <https://doi.org/10.4232/1.8013>.
- Google. 2024. *Google Search*. <https://www.google.com>.
- GTD. 2020. *START (National Consortium for the Study of Terrorism and Responses to Terrorism)*. (2020). *Global Terrorism Database 1970 - 2020*. <https://www.start.umd.edu/gtd>.
- Karapanagiotis, Pantelis, and Marius Liebald. 2023. "Entity Matching with Similarity Encoding: A Supervised Learning Recommendation Framework for Linking (Big) Data." *SAFE Working Paper No. 398*, <https://doi.org/10.2139/ssrn.4541376>.
- LexisNexis. 2025. *LexisNexis Database*. <https://www.lexisnexis.com>.
- Remus, Robert, Uwe Quasthoff, and Gerhard Heyer. 2010. "SentiWS-A Publicly Available German-language Resource for Sentiment Analysis." In *LREC*.