

DISCUSSION PAPER SERIES

DP14455

**TERRORIST ATTACKS, CULTURAL
INCIDENTS AND THE VOTE FOR
RADICAL PARTIES: ANALYZING TEXT
FROM TWITTER**

Francesco Giavazzi, Giacomo Lemoli, Gaia Rubera
and Felix Iglhaut

**MACROECONOMICS AND GROWTH
PUBLIC ECONOMICS**



TERRORIST ATTACKS, CULTURAL INCIDENTS AND THE VOTE FOR RADICAL PARTIES: ANALYZING TEXT FROM TWITTER

Francesco Giavazzi, Giacomo Lemoli, Gaia Rubera and Felix Iglhaut

Discussion Paper DP14455
Published 29 February 2020
Submitted 27 February 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Macroeconomics and Growth
- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Francesco Giavazzi, Giacomo Lemoli, Gaia Rubera and Felix Iglhaut

TERRORIST ATTACKS, CULTURAL INCIDENTS AND THE VOTE FOR RADICAL PARTIES: ANALYZING TEXT FROM TWITTER

Abstract

We study the role of perceived threats from cultural diversity induced by terrorist attacks and a salient criminal event on public discourse and voters' support for far-right parties. We first develop a rule which allocates Twitter users in Germany to electoral districts and then use a machine learning method to compute measures of textual similarity between the tweets they produce and tweets by accounts of the main German parties. Using the dates of the aforementioned exogenous events we estimate constituency-level shifts in similarity to party language. We find that following these events Twitter text becomes on average more similar to that of the main far-right party, AfD, while the opposite happens for some of the other parties. Regressing estimated shifts in similarity on changes in vote shares between federal elections we find a significant association. Our results point to the role of perceived threats on the success of nationalist parties.

JEL Classification: C45, D72, H56

Keywords: National elections, Political parties, Text Analysis, social media, Terrorism

Francesco Giavazzi - francesco.giavazzi@unibocconi.it
Bocconi University, IGIER and CEPR

Giacomo Lemoli - gl1759@nyu.edu
New York University

Gaia Rubera - gaia.rubera@unibocconi.it
Bocconi University

Felix Iglhaut - felix.iglhaut@unibocconi.it
Bocconi University, IGIER

Acknowledgements

We thank Guido Tabellini, Marco Tabellini, Jesse Shapiro and participants in the Spring 2019 NBER Political Economy meeting for very helpful comments. We gratefully acknowledge a grant from the Italian Ministry for Universities MIUR PRIN Prot. 2015FMRE5X.

Terrorist Attacks, Cultural Incidents and the Vote for Radical Parties: Analyzing Text from Twitter*

Francesco Giavazzi[†], Felix Iglhaut[†], Giacomo Lemoli[‡] and Gaia Rubera[§]

February 27, 2020

Abstract

We study the role of perceived threats from cultural diversity induced by terrorist attacks and a salient criminal event on public discourse and voters' support for far-right parties. We first develop a rule which allocates Twitter users in Germany to electoral districts and then use a machine learning method to compute measures of textual similarity between the tweets they produce and tweets by accounts of the main German parties. Using the dates of the aforementioned exogenous events we estimate constituency-level shifts in similarity to party language. We find that following these events Twitter text becomes on average more similar to that of the main far-right party, AfD, while the opposite happens for some of the other parties. Regressing estimated shifts in similarity on changes in vote shares between federal elections we find a significant association. Our results point to the role of perceived threats on the success of nationalist parties.

*We thank Guido Tabellini, Marco Tabellini, Jesse Shapiro and participants in the Spring 2019 NBER Political Economy meeting for very helpful comments. We gratefully acknowledge a grant from the Italian Ministry for Universities MIUR PRIN Prot. 2015FMRE5X.

[†]Bocconi University and IGIER

[‡]New York University

[§]Bocconi University

In the past decade economic hardship following major global shocks and the increasing salience of market openness and immigration in the public discourse have accompanied the rise of protectionist and culturally conservative politicians generally opposed to globalization understood as the free circulation of goods and people. Meanwhile the vote share of nationalist, far-right, often racist parties also increased in many Western democracies. The tight connection of economics and culture in these political platforms has led scholars and commentators to discuss which of the two is more important in explaining the anti-globalization backlash of Western voters.

At the same time, Europe faced in the second half of 2010s an unprecedented sequence of religiously motivated terrorist attacks, accompanied (though not caused) by an intensification of migration flows, which has made the defense of national borders an even more salient political issue. Far-right parties have framed some of their anti-immigration stances as policies designed to provide security against the threat posed by foreigners (Couttenier et al. 2019).¹

Disentangling the role of economics and culture is not easy, since in recent years economic distress and migration flows have happened in parallel in several Western democracies, as some of the European countries more geographically exposed to the refugee crisis of the 2010s were still struggling with the consequences of the financial crisis.

In this paper we investigate the extent to which perceived threats associated with terrorist attacks have influenced public opinion and the support of far-right parties. We use exogenous time variation in terrorist attacks and a salient crime event to disentangle the roles of culture from economics.

We study the case of Germany which, despite a fast recovery from the Great Recession and a low unemployment rate, has experienced the rapid rise of far-right party, AfD (*Alternative für Deutschland*). In the period running from the 2013 to the 2017 Federal elections (Bundestag elections), some of the terrorist attacks among the ones perpetrated by Islamist militants in Western Europe occurred in Germany. Moreover, in the midst of the refugee crisis, some criminal acts perpetrated in Germany by men of reported Arab and Middle Eastern origin became very salient. These events fueled the political debate over the consequences of the government's immigration policies. AfD politicians have identified the open border policy as posing a security threat for the German population.²

At the end of this period, in the 2017 Federal elections, the AfD entered the lower house of Parliament by almost tripling its vote share, a step no far-right party has accomplished since the immediate post-World War II history. We surmise that events increasing the perception of threats

¹For instance, Matteo Salvini, then right-wing leader in Italy, said "the risk of terrorism is incredibly high [...] we ask for a tight control of all our borders and the suspension of any further landing on our coasts" (Corriere Della Sera 28.03.2018).

²Among many others, the former leader of the party, Alexander Gauland, openly advocated the closing of German borders by all means, even to the point of tolerating the cruel pictures which might result from it, in order to not be blackmailed by the eyes of children (Zeit Online 24.02.2017).

from other cultures, at a time when such threats were very visible and salient, may have shifted voters' attitudes towards those affirmed by the party most vehemently opposed to immigration, and eventually may have had an effect on electoral outcomes.

In order to provide empirical evidence for this hypothesis we use Twitter data. First, we download the tweets posted by the official national Twitter accounts of the seven main German parties and perform some basic analysis revealing the political leaning of those messages. We then geolocalize a sample of more than 178,000 Twitter users and collect all their available tweets to obtain an electoral constituency panel dataset at daily frequency. Bringing together parties and users, we use a Natural Language Processing algorithm (`doc2vec`) to compute a daily measure of similarity between the content posted by parties and the content posted by Twitter users in a given constituency. We use this measure of similarity to infer the alignment of Twitter users with national parties.

Then, we identify a set of terrorist attacks and a criminal event related to cultural diversity. We use time variation in text similarity and the exogenous timing of these events to estimate a discontinuous growth model (Bliese et al. 2016). This allows comparing the predicted similarity in the presence and in the absence of these events.

We find that following these events the tweets posted in German constituencies become, on average, more similar to the AfD's tweets and less similar to the other parties' tweets. We could attribute these results to the AfD strategically changing its language to make it more similar to that used by the public. However, we conduct a within-party analysis of tweets over time and find no evidence that party accounts change their language in the aftermath of our events. Thus, we can rule out the possibility that a strategic change in the AfD language drives our findings. Hence, it seems plausible that the increasing similarity between German Twitter users' and AfD' language is driven by Twitter users changing theirs to become more similar to the AfD language, including a negative connotation of diversity and immigration.

We further find that standard economic variables do not explain the estimated change in language similarity after an event. These shifts, however, are significantly correlated with the difference in vote shares obtained by parties between the two elections.

A few contributions stem from our results. First, we speak to the literature on the roots of far-right support (Colantone et al. 2018c; Colantone et al. 2018a; Ballard-Rosa et al. 2018; Inglehart et al. 2016) by emphasizing the role of perceived threats from other cultures which arise as a consequence of terrorist and crime events.

We also contribute to the literature on the effects of terrorism and crime on public opinion (Echebarria-Echabe et al. 2006; Finseraas et al. 2011; Legewie 2013; Ferrín et al. 2019). We advance this literature by measuring people's attitudes through social media text and showing

that they become closer to the policy views of national parties. Twitter data allow us to exploit high-frequency information and to avoid relying on large surveys (Curini et al. 2015). We finally contribute to the methodology of social science research based on Twitter data, by providing a strategy to geo-locate Twitter users to geographic units using following patterns and textual similarity to measure people’s attitudes and party preferences.

The structure of the paper is the following: we first discuss the literature related to our paper and present the historical background of our study. Then we describe the data and some descriptives from our body of tweets, we introduce our measurement and empirical strategy, and we present and discuss our main results before concluding.

Related Literature

Economic and cultural shocks

A large and growing literature has studied the determinants of the political backlash against globalization which has favored the rise of nationalist and isolationist parties. An influential strand of this research has focused on the political consequences of job displacements induced by global trade, which have undermined the system of embedded liberalism adopted by most traditional parties (Colantone et al. 2018a; Colantone et al. 2018c). Moreover, several studies document a tight connection between the economic costs of globalization and worsening attitudes towards immigration and cultural diversity (Colantone et al. 2018b; Bisbee 2019; Cerrato et al. 2018) and towards democratic values and institutions (Colantone et al. 2018b).

Another influential strand of literature focuses on the effects of globalization on voting through induced changes of social identity. The electoral support for Donald Trump is partly associated with the decline of a traditional way-of-life based on the relative status enjoyed by some social groups (Mutz 2018; Ballard-Rosa et al. 2018; Jardina 2019; Baccini et al. 2019). Other work posit, more generally, that the success of nationalist politicians rests on the reaction of social groups holding traditional sets of values against cosmopolitanism, multiculturalism, and diversity (Inglehart et al. 2016). Opposition to immigration is a crucial component of this backlash. Indeed, several studies identify a causal connection between immigration and refugee settlement and shifts in voting preferences towards the right (Barone et al. 2016; Halla et al. 2017; Dustmann et al. forthcoming), although with some exceptions (Steinmayr 2018; Hill et al. 2019). This literature generally finds that non-economic factors such as concerns about inter-ethnic contact and deterioration of local public goods explain a large part of the change in voters’ preferences.

Hangartner et al. (2017) find that the mere *exposure* to immigrants, without expectation of labor market costs, may increase native hostility towards immigrants and minorities. Tabellini

(forthcoming) finds in the American context that immigration generated substantial political opposition even in cases where it was beneficial for the economy, particularly when the cultural distance between natives and newcomers was high.

The political consequences of terrorism

Terrorism and crime may have a significant impact on public opinion and electoral outcomes. In Western democracies, when attacks have been religiously motivated, the political debate that ensued has been largely concerned with issues of immigration policy and government attitudes towards ethnic and religious minorities. The 9/11 attacks impacted negatively on attitudes against Muslims (Schüller 2016) and positively on the number of anti-Muslim hate crimes in the U.S. (Gould et al. 2016), and are associated with higher discrimination against other minorities at large (McConnell et al. 2018). Similar effects have been observed in Europe in the aftermath of relevant terrorist events (Legewie 2013; Echebarria-Echabe et al. 2006; Ferrín et al. 2019). The effects of terrorism and of its threat often persist in the ballot box, increasing support for right-wing parties (Getmansky et al. 2014; Kibris 2011) and polarization (Berrebi et al. 2008).

Parties, voters, and social media

With social media becoming increasingly important platforms for political communication and information, online interactions between users and politicians have become more sophisticated. If in the past doubts have been cast on the efficacy of social media for spreading political information in all contexts (Vaccari et al. 2013), now politicians often follow online discussions in the choice of relevant topics (Barberá et al. 2019).

Political views expressed on social media are also connected to offline engagement: Twitter users engaging in discussions about politics appear to be more likely to interact directly with political actors in their daily life (Vaccari et al. 2015). Social media are also effective as drivers and coordination devices of participation in protests (Jost et al. 2018; Larson et al. 2019). Existing research has assessed political partisanship of Twitter users from their follower networks (Barberá 2015; Eady et al. 2019) and from textual features, as language is related to partisan alignment (Gentzkow et al. 2010; Jones et al. 2018). Estimates of political alignment have then been used to study online behavior (Eady et al. 2019). Twitter, like other social media (Müller et al. 2019), is also a relevant platform for the diffusion of anti-minority hatred (Müller et al. 2020).

From a methodological point of view, an appealing feature of Twitter is that, differently from surveys, it allows to identify variation in attitudes in real time (Curini et al. 2015) and around specific events. This makes the analysis of tweet language promising to study shifts of opinions and sentiment in settings where surveys are unavailable or the relevant issues are potentially sensitive

and subject to social desirability bias.

Background

The rise of the far-right in German politics

AfD was founded in 2013 as a fiscally conservative party defending German public finances in the context of the European Union programs, monetary union in particular. The party ran in the 2013 Federal elections and missed the 5% threshold of nationwide votes for entering the German Parliament.³ At the European elections in the following year, AfD gained 7.1% of nation-wide votes and entered the European Parliament.

In July 2015, following a change in leadership, a new chair was appointed, Frauke Petry, and most of the moderate members left. From that moment on, the party agenda shifted from primarily economic issues to strong criticism towards immigration policy and an emphasis on nationalism and social conservatism. The communication of AfD on social media also changed, with words such as “migration” and “Islam” trending upwards in the posts (Cantoni et al. 2019).

Between 2015 and 2017, AfD performed well in all state elections and at the 2017 Federal elections it received the third largest number of votes for the party list and was the third largest to be represented in the *Bundestag*. The electoral success further continued in state elections until 2019.

Electoral support for AfD is correlated with concerns about immigration policy and distrust against political institutions (Hansen et al. 2019) and has a strong regional component, with a solid constituency in eastern states. The party also uses actively social media to spread political content and influence the debate (Darius et al. 2019).

The Refugee Crisis and Terrorism

The AfD rise coincided with the dramatic increase in migration and refugee flows to European countries, including Germany. Between 2014 and 2017, almost 4 million people applied for asylum in some European country (Eurostat 2019). As a reaction to the worsening of the civil conflict in Syria, in 2015 Chancellor Angela Merkel announced that the government would allow Syrian asylum-seekers to cross the German borders. The refugee policy soon became one of the crucial issues around which AfD concentrated its opposition to the government led by Angela Merkel.

Parallel to the refugee crisis, Western Europe was hit by an unprecedented wave of fundamentalist Islamic terrorism. Starting with the shooting in the headquarter of magazine *Charlie Hebdo*

³We refer to the votes for party lists. In the German electoral system voters also cast a vote for a constituency representative, elected by relative majority.

in Paris (January 2015), attacks were perpetrated in France, Belgium, the United Kingdom, and Germany. International and domestic attacks fueled criticisms of Merkel’s government and its immigration policy.

Political debate in Germany saw immigration policy becoming even more salient when, in the aftermath of 2016 New Year’s Day, German media reported that, during Year’s Eve, groups of men, described by some witnesses as mostly of “Arab appearance”, sexually assaulted hundreds of women in several German cities (Zeit Online 05.01.2016). The scale of the event and the identification of perpetrators with the Muslim religion sparked public outrage. As a result, tensions over security and multiculturalism characterized the political climate before the 2017 federal election.

Data

Parties

We analyze the parties that won seats in the federal parliament (*Bundestag*) in 2017: *Alternative für Deutschland* (AfD, Alternative for Germany), *BÜNDNIS 90/DIE GRÜNEN* (The Greens), *Christlich Demokratische Union Deutschlands* (CDU, Christian Democratic Union for Germany), *Christlich-Soziale Union in Bayern* (CSU, Christian Social Union in Bavaria), *Die Linke* (The Left), *Freie Demokratische Partei* (FDP, Free Democratic Party), and *Sozialdemokratische Partei Deutschlands* (SPD, Social Democratic Party of Germany). For each party, we consider the main, national-level Twitter account.⁴

Electoral and Structural Data

For the years of federal elections, the Federal Returning Officer of Germany (Der Bundeswahlleiter 2017) publishes the election results of each electoral constituency. We use the vote for the party list, which it is the most important for the distribution of seats in the German Bundestag, rather than preferences for individual candidates.

For each constituency, the Federal Statistical Office (FSO) also publishes a set of aggregate structural variables, such as standard economic variables and demographic variables. Since electoral constituencies do not follow the borders of (NUTS-3) administrative districts, the FSO publishes these statistics for the federal election years only. From the FSO we also collect data on the number of refugees in each electoral constituency at the administrative district level and aggregate them at the constituency level.

⁴We exclude the party leaders’ and representatives’ personal accounts in order to assess comparable accounts for all parties. This is because some party leaders, such as Angela Merkel, are not on Twitter.

Twitter Users

We construct a sample of German Twitter users, which encompasses most of German electoral constituencies. We start from a complete list of towns belonging to each constituency provided by the Federal Returning Officer. The first challenge is to identify where Twitter users live, i.e. the town where they are registered to vote. Twitter users can voluntarily choose to publish any location they wish on their profile and there is no reliable way to double check the provided information. Hence, using information provided by users would lead to four possible outcomes: missing addresses, reported correct addresses, reported incorrect addresses, and reported fantasy addresses (such as Disneyland). Excluding the latter is straightforward, but there is no simple method to verify whether the location an user provides is her real place of residency or not. For this reason, we construct a rule that allocates users to a constituency, whether or not they provide information on their location.




Figure 1: Sampling Rule

The rule works as follows. The 299 German electoral constituencies⁵ are drawn with the goal of equalizing population across them. Thus, electoral borders in general do not follow a common structure, but are drawn over towns and districts. By the end of 2017 there existed 401 districts and district-free cities⁶, which correspond to the NUTS-3 classification of the European Union. For a given constituency, our approach first identifies the largest towns within each district of the given constituency. Here we face two possible situations (as shown in Figure 1).

For a constituency (such as C_1 , one of the two squares in Figure 1) that contains parts of one or more districts (D_1, D_2) we consider the largest towns in the respective districts belonging to the constituency (here, T_1 as the largest town of district D_1 within constituency C_1 , and T_2 as the largest town of district D_2 within constituency C_1). Because one district can overlap with several constituencies (here D_2 is part of C_1 , C_2 as well as other non-labeled area), the chosen towns are not necessarily the largest towns in their districts (T_3 and T_2 both belong to district D_2 , and T_3

⁵This number refers to constituencies for the general election of 2017.

⁶District-free cities are of considerable population size to have their own administration, while cities and towns belonging to districts share parts of the administration.

might be larger than T_2 . Nevertheless T_2 is the largest town within D_2 that is still part of C_1 . The second case concerns multiple constituencies (C_3 to C_6) which are entirely located within a district-free city (T_5). For instance, the city of Berlin is divided into eleven constituencies. In these cases we merge all constituencies into one. Our final sample comprises 261 constituencies, either original or artificially merged, in which the rule described above produces a sample of 493 towns. For constituencies belonging to case 1 (Figure 1a), our rule usually includes two or three towns, depending on how many districts intersect a constituency.

Within those towns we identify the Twitter accounts of so called "landmarks". These are public or commercial accounts which can be clearly located in a given town and are likely to be followed by residents. Examples are small-scale shops, town halls, police stations, fire departments or, theaters. We do not consider sport clubs, TV stations or newspapers as landmarks, because their attraction is not bound to local residents. For example, following a famous football club or a well-established newspaper is not a reliable source to infer where a user lives. Similarly, the catchment area of possible landmarks in constituencies outside of towns is much less clear than for landmarks within a town. For example, large shopping centers might attract people from relatively far away towns and using them can lead to wrong attributions to a constituency. This strategy produced a sample of 5,231 landmark Twitter accounts, around ten per town in our sample.

Having identified local landmarks, we use the Twitter API to fetch all the followers of those landmarks. We eliminate those users who follow less than three landmarks in the same constituency or follow landmarks in more than one constituency: that is we assume that people who follow at least three landmarks of a certain constituency and no landmarks of another constituency live there. This strategy produce a sample of 178,271 located Twitter users. This sampling procedure has the advantage of limiting the risk of including non-human users (bots) in our sample, which instead may dramatically influence the political debate on social networks. Indeed, bots are very unlikely to follow accounts of facilities at a very local level, such as our landmarks (Ferrara et al. 2016).

For the users in our sample, we download all available tweets. Twitter limits the access to roughly the latest 3,200 tweets, but since only 128 users in the sample tweeted more than this, we consider the influence of this limit negligible and conclude that we use essentially all the tweets that the users in our sample posted.

We also perform a first basic electoral analysis using a pre-existing corpus of about 20,000 German tweets by far-right supporters containing racist or violent rhetoric. This corpus of "hate speech" is a subset of The Polly Corpus (De Smedt et al. 2018), a multimodal study corpus of online political debate in Germany, which includes tweets about politicians, by politicians, their fans and far-right supporters.

Possible Sources of Bias

Three possible sources of bias could be present in our data.

First, we can retrieve Twitter users in only 225 constituencies out of the 261. This is due to the fact that for some constituencies we could not geo-localize a sufficient large number of users. Bias might arise if the constituencies in our sample were more supportive of the AfD than those that we do not observe. However, we find that the AfD vote share is slightly lower in the constituencies that we observe than in those we do not, which suggests that our sample does not consist of above-average right-wing supporters (Table 1).

Table 1: AfD Vote: In-sample vs Out-of-sample

	In-sample	Out-of-sample	Full Sample
AfD Vote 2017 (Mean)	0.1280	0.1466	0.1306
Δ AfD Vote (Mean)	0.0809	0.0981	0.0832
N	225	36	261

Notes: Δ AfD Vote refers to the difference in vote share from 2013 to 2017.

The second possible source of bias is due to the fact that, within the constituencies that we observe, we have more landmarks, and hence more Twitter users, in large cities than in smaller towns. It is no surprise that we find more landmarks in larger cities, as here there are naturally more facilities which qualify as a landmark. This sampling issue would bias our results if users in larger cities would support AfD more than users in smaller cities. However, since support for the AfD is highest in rural areas with low population density, we believe that the bias would likely be against the inference of a non-zero effect and therefore our estimates should represent a lower bound.

Finally, the third source of bias comes from the fact that a Twitter user might be different from a representative German voter. The exact amount of active German Twitter users is unknown; different sources estimate it between 2 and 5 million users over a population of about 83 million.⁷ There is clearly a self-selection mechanism in our sample. However, two pieces of evidence suggest that, also in this case, our sample does not consist of an over-proportional number of right-wing supporters. First, based on the representative electoral statistics for 2017, we know that 66% of AfD voters are more than 44 years old (Kobold et al. 2018).⁸ A survey conducted in 2017/18 reports that Twitter was used by 5-9% of the total population of Germans younger than 50 years. In contrast, less than 2% of Germans being 50 years or older were registered on the social network

⁷In the United States this figure is about three times larger.

⁸The representative electoral statistics are not a survey. These statistics are constructed based on a sample of official ballot papers indicating the true gender and age group of a voter before the vote is cast. Detailed information can be obtained from The Federal Returning Officer.

(Frees et al. 2018). Hence, we conclude that while older people are over-represented among AfD supporters, younger people are over-represented among German Twitter users.

Table 2: Twitter Accounts of Major German Parties

Party	Party Account	# tweets	# Followers	Joined
AFD	@AfD	18,600	130,000	Sep-12
Bündnis 90/ Die Grünen	@Die_Gruenen	18,000	441,000	Apr-08
CDU	@CDU	16,300	274,000	Feb-09
CSU	@CSU	14,800	186,000	Feb-09
Die Linke	@dieLinke	24,500	254,000	Jun-09
FDP	@fdp	10,900	331,000	May-09
SPD	@spdde	32,200	354,000	Mar-09

Notes: Retrieved February 11, 2019. Amount of tweets includes retweets.

Of course, this evidence does not reject a right-wing bias in our sample. The AfD still has supporters of young age and there is still a chance that a large share of Twitter users belong to this group. However, Table 2 clearly shows that, based on the pattern of followers, the AfD is not as popular as other parties on Twitter. For instance, while the left party The Greens posted roughly the same amount of tweets and retweets (although over a longer period of time), it has more than three times as many followers, compared to the number of AfD followers. In fact, the AfD is the party with the fewest followers on Twitter, although it exceeded three of those parties in vote share. We believe that this represents strong evidence that Twitter users are not overly supportive of the AfD.

The bottom line is that the demographic profile of Twitter users in Germany is not close to that of a representative supporter of a right-wing populist party. Hence, a representative Twitter user in our sample is likely to be more moderate or liberal in political beliefs than a potential AfD voter. Therefore, we surmise that although we will not be able to identify an effect for the German electorate, our method will most likely underestimate it. If there is an effect in the population of Twitter users, there should be an even stronger effect in the German population.

Events

For our analysis we use eleven events, from the end of 2015 until close to the federal election in 2017. We choose these events because they represent large shocks to public opinion. Among the several events related to terrorism and crime reported in the media between 2015 and 2017, we look for a subset which satisfies some properties. First, they need to be plausibly exogenous

to local conditions. Hence, we disregard very local incidents such as thefts or small-scale violent incidents. Events happening in other countries are particularly appropriate to this goal. Second, they need to be large shocks, affecting public opinion not only in the area where they took place (i.e. town or district), but in the whole country. Thus, we exclude some non-deadly attacks and relatively less important events. Third, we select events that plausibly highlight the salience of an external cultural threat: since all the attacks in this period had Islamist extremism as alleged or clear motivation, we believe this presumption is realistic. These are the ten terrorist attacks that we consider:

Table 3: Terrorist events

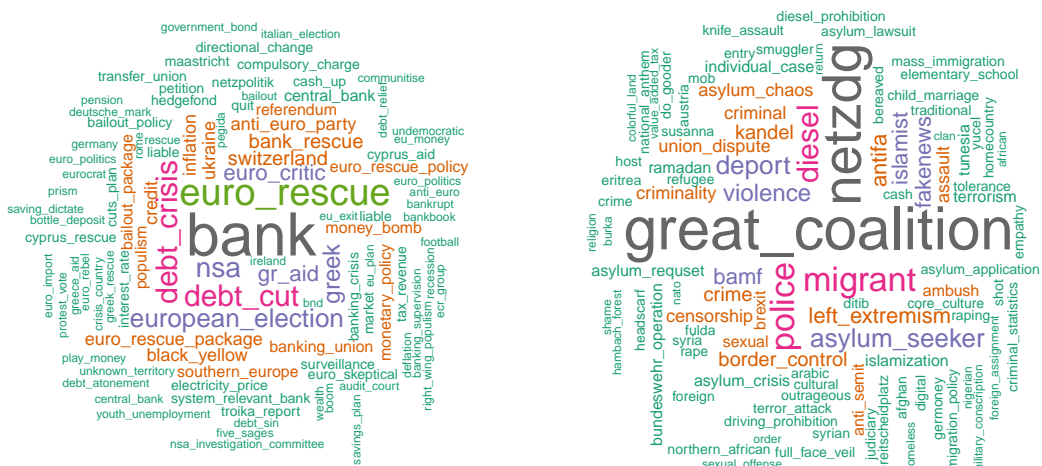
Date	City	Circumstance	Fatalities
November 13, 2015	Paris, France	Simultaneous attacks by groups of terrorists on several targets, including the <i>Bataclan</i> concert hall.	130
March 22, 2016	Brussels, Belgium	Coordinated bombings at several locations.	32
July 14, 2016	Nice, France	Truck driven at high speed over the crowd.	86
December 19, 2016	Berlin, Germany	Truck driven over the crowd in a Christmas market.	12
March 22, 2017	London, UK	Car driven over pedestrians.	5
April 20, 2017	Paris, France	Three policemen and another person shot by an attacker.	3
May 22, 2017	Manchester, UK	Suicide bombing after a concert at Manchester Arena.	22
June 3, 2017	London, UK	Car driven over pedestrians.	8
August 16 2017	Barcelona, Spain	Bombs detonated and a car driven over pedestrians.	16
September 15, 2017	London, UK	Bomb detonated at a train station.	0 (30 injured)

In addition, we include a non-terrorist event which shocked public opinion in Germany and across Europe and generated wide political and social reactions consistent with the idea of cultural threat. In *December 31, 2015 and January 1, 2016 in Cologne, Germany*, during the New Year’s Eve celebrations, several hundred women were subject to mass harassment and sexual assaults. According to the police, investigations on the perpetrators concentrated on North African and Syrian young men. Similar cases were later reported in other German cities.

In order to ease notation, we will from now on use the term *events* referring both to terrorist attacks and the non-terrorist crime incident just described.

Tweets and Content

Before proceeding with our main analysis of the eleven events just described, we provide some information about the tweets we collected. For political parties, if the language used on Twitter is representative of the party language, we would expect to see strong differences in language across



(a) AfD: Before July 2015 (b) AfD: After July 2015

Figure 2: Within Party Comparison: AfD before and after July 2015

very opposite parties, and within a party across time in case a party substantially changes its position. Furthermore, as we are able to locate Twitter users within constituencies, we can analyze correlations between the language used in each constituency and electoral results.

Parties’ Tweets: Transformation of the AfD

As our main interest is focused on the AfD, we present evidence for the transformation of the party by analyzing how the language of the AfD changed over time. As described above, in 2015 the AfD took a remarkable turn from a fiscally conservative euro-skeptic party to an outright far-right party. The crucial date for this turn is July 2015. Figure 2 shows two word clouds of the words that the party was most likely to use before and after this date, respectively. The technical details on how this word cloud was created are in Appendix A. In short, taken two documents (here, AfD before and after 2015), we compute the log-odds-ratios for all the words in these documents, thus identifying which words are most likely to appear in tweets before July 2015, and at the same time least likely to appear after that date. This allows to identify the topics contrasting the two periods (here, AfD before and after 2015) pointing out. As expected, for a categorized list of politically relevant words we see a clear separation of topics.⁹ The picture on the left is clearly dominated by economic-related phrases, concerning the European debt crisis and monetary policy. On the right, the overall image is characterized by socially conservative words, mainly related to crime (*sexual_offense*, *rape*, *violence*), Islam (*ramadan*, *islamist*, *burka*, *headscarf*), immigration policies and refugees.

⁹We excluded words without meaning or political contexts and those which were incomprehensible after the pre-processing from the wordclouds. See Appendix A for details.

Leaving the sphere of a single party, we would expect to see strong differences in the language of very opposite parties. When comparing the right-wing AfD with the left-wing party The Greens over the whole time interval for which we collected tweets, AfD reveals a mix of economic topics, mainly concerned with the European debt crisis, and topics related to immigration as well as clear right-wing rhetoric. On the other hand, the picture for The Greens shows the expected focus on environmental topics mixed with traditional leftist topics. The resulting wordcloud is shown in Figure A.1 in Appendix A. This evidence provides initial support to our choice of using tweets language to proxy for party positions.

Constituency Tweets: Racist Words and Votes

We use a corpus of tweets containing hate speech provided by De Smedt et al. (2018) to manually extract a list of right-wing leaning or racist words.¹⁰ Given that we are able to locate Twitter users inside German electoral constituencies, we would expect, within a constituency, a positive correlation between the share of right-leaning words (out of the total amount of words) and the vote for a right-wing party. This is exactly what we find. The share of racist words is positively and significantly correlated with the change in votes for the AfD. Furthermore, we find no significant correlation with other parties except for the most left party (The Left), which shows a negative significant correlation. Results are presented in Table A.1 in Appendix A.

This finding suggests that our sample of German tweets offers possibilities for political analysis. Also, it corroborates the validity of our geo-localization procedure. We calculated the frequency of racist words inside an electoral constituency using the language of Twitter users located inside the constituency. If the geo-localization procedure performed poorly, a correlation between the frequency of racist words and the vote of a right-wing party would be unlikely.

Overall, the analysis of how parties' tweets change over time and the analysis just described corroborate our choice of using our Twitter textual data to investigate the role of events in moving German Twitter users closer or farther away from political parties.

Similarity between Texts

We compute a daily similarity between the tweets of the parties and the tweets of each constituency by transforming the two groups of tweets into vectors with doc2vec, a deep learning technique. Details on prior steps regarding text preprocessing, the functioning of doc2vec, and the hyperparameters used are provided in Appendix B. In the following, we briefly summarize the method.

¹⁰As there are no clear rules defining what makes a word racist, the list we extract contains only words of extreme political language, often dealing with insults to ethnic groups or references to the Third Reich. We provided our list to a native German independent of this project and in case of doubt, a word was excluded from further usage.

For our analysis, we create two “documents” at daily frequency: a party- and a constituency-document. The party document is the text of all the tweets the party posted on a certain day. A constituency document is the text of all the tweets that all the users in our sample located in a given constituency posted on a certain day. Since we have 752 days in our observation period (from September 4th 2015 to September 24th, 2017¹¹), we end up with 752 documents for each party and 752 documents for each constituency in our sample.¹²

Given documents containing either tweets of a party or tweets of a constituency on a given day, we use doc2vec (Le et al. 2014), an unsupervised deep learning algorithm that learns how to represent each document with a unique vector. We then measure similarity between party p and constituency c in day t as the cosine similarity between the two corresponding vectors:

$$\cos \theta_{cp_t} = \frac{\vec{c}_t \cdot \vec{p}_t}{\|\vec{c}_t\| \|\vec{p}_t\|}$$

This is the measure of daily similarity between each constituency and each party used in our empirical analysis.

Validation

We perform three validation exercises to control whether cosine similarity is a valid measurement for how close public opinion is to the national parties. The details of this validation are in Appendix C, but we provide here the intuition and the basic results. As a first check, for a sub-sample of tweets we had a human reader categorize each tweet as similar or dissimilar from those posted by different parties on the same day. We find the human evaluation to be qualitatively similar to the one estimated by the algorithm. Second, we compute similarity of party tweets with a sample of tweets whose content is known *a priori*: using the corpus of German tweets containing racist words discussed above, we find that the AfD is indeed the party whose language is the closest to this racist corpus for the whole period of time, and even more so for certain sub-periods. Third, we correlate our similarity measure with electoral results at the national level and regular polling data at the state level and find that these measures are positively and significantly associated.

¹¹As stated above, July 2015 marked a turning point in the history of the AfD. We leave two months between the change in leadership of the AfD and the starting point of our analysis, but the empirical method chosen is not sensitive to the exact day. More importantly, we choose this day as it represents a major peak of events during the 2015 refugee crisis with the German government opening borders for refugees coming from Hungary.

¹²752 is the maximum possible amount of documents for a given constituency in case the users posted tweets every single day.

Empirical Strategy

Effect of Events

We are interested in identifying the association of the set of events with the support for political parties at the electoral constituency level in the following Federal election. Our analysis relies on the plausible assumption that this subset of events represents exogenous shocks to public opinion whose occurrence is independent of any local conditions. Hence, we use this exogenous time variation in events to disentangle the roles of culture from economics and other local factors. The size of the possible effect of a specific event, however, could differ across constituencies because of their different characteristics. In other words, the degree to which a population reacts to an event is unlikely to be uniform.

We measure similarity to all parties across German electoral constituencies and over time, resulting in a panel data structure with daily frequency. One way to study the effect of events on similarity is to compute the difference between the similarity prior to an event and the one after it happened. However, inference based on this value has drawbacks. First, there could be self-selection into tweeting: that is, people who use Twitter to comment terrorist attacks while they happen, or minutes after, may not be representative of the overall Twitter population of that constituency. Moreover, we could simply measure an immediate outrage, while what we are interested in is the long-run effect of those events on the similarity between people and parties. That is, we want to investigate whether there exists a positive or negative shift towards parties that occurs at the time of those events.

To measure this long-term shift in similarity we use a discontinuous growth model (Bliese et al. 2016). This model is able to examine the evolution over time of a time series, such as similarity, punctuated by one or more discontinuities. Consider Figure 3, a simple visualization of the model. It allows at specified points in time for a change in growth (slope) and level (intercept) of the time series of interest. In our case, after each event, both the time trend and the level of similarity to parties are allowed to shift. The change in trend and level is relative to a trend in the absence of any discontinuity, and is estimated based on the similarity since the last event. In other words, the coefficients of change are not based on a specific day, thus overcoming the issue of self-selecting into tweeting after events. The discontinuous growth model thus enable us to move beyond the immediate reaction to an event and capture its long-term effect on the evolution of similarity over time (Bliese et al. 2016).

To account for the possibility that reactions to events differ across constituencies, we estimate a model with random coefficients at the constituency-level. To check whether there is variation across constituencies, we run an unconditional means model with random coefficients and estimate




Figure 3: Discontinuous Growth Model: Simple Visualization

the proportion of total variance that occurs between constituencies. Since we find it to be about ten percent (results available upon request), we argue that there is enough variance to justify the use of random coefficients (Singer 1998). Omitting these random coefficients would lead to biased estimates and standard errors (Goldstein 2013). We thus allow for changes in intercept and time trend of similarity to vary across constituencies on the day of each event. Given the eleven events, for each party p we estimate separately the discontinuous growth model using maximum likelihood

$$\begin{aligned}
 \text{simil}_{ti}^p &= \pi_{0i}^p + \pi_{1i}^p \text{Time}_t + \pi_{2i}^p \text{Time}_t^2 + \pi_{3i}^p \text{Year} \\
 &+ \sum_{k=1}^{11} \left[\pi_{4,ki}^p E_{kt} + \pi_{5,ki}^p \text{Reset}_{kt} + \pi_{6,k}^p \text{Reset}_{kt}^2 \right] + \epsilon_{ti}^p
 \end{aligned} \tag{1}$$

where a coefficient with subscript i consists of a fixed and a random component, that is

$$\begin{aligned}
 \pi_{0i}^p &= \pi_0^p + r_{0i}^p, \\
 \pi_{1i}^p &= \pi_1^p + r_{1i}^p, \\
 \pi_{4,ki}^p &= \pi_{4,k}^p + r_{4,ki}^p \quad \forall k \in \{1, \dots, 11\}, \\
 \pi_{5,ki}^p &= \pi_{5,k}^p + r_{5,ki}^p \quad \forall k \in \{1, \dots, 11\}
 \end{aligned} \tag{2}$$

and error terms and random coefficients are independently distributed as

$$\epsilon_{ti}^p \sim N(0, \sigma_p^2), \quad \mathbf{r}_i^p \sim N(\mathbf{0}, \Sigma_p), \quad \epsilon_{ti}^p \perp \mathbf{r}_i^p$$

Σ_p is a diagonal matrix, that is random coefficients are assumed to be jointly independently and identically distributed.

simil_{ti}^p is the measured daily similarity for party p in constituency i over all periods; Time_t and Time_t^2 are a time and a quadratic time trend: their coefficients estimate how similarity would

evolve in the absence of events¹³; r_{1i}^p are the random coefficients allowing for between-constituencies differences in time trend¹⁴; $Year$ is a dummy equal to 1 in 2016 and 0 elsewhere.¹⁵ For $k = 1, \dots, 11$ E_{kt} is the event k indicator variable, coded 1 after an event has occurred until the next event occurred, and 0 otherwise: the associated parameter $\pi_{4ki}^p = \pi_{4k}^p + r_{4ki}^p$ estimates the extent to which the predicted value of this model on the day of event k differs from the predicted value in absence of *any* event, and is based on the trend prior to the first event. In other words, we are estimating the difference between predicted similarity after events and the predicted counterfactual in the absence of any event. $Reset_{kt}$ and $Reset_{kt}^2$ are event-specific variables coded 0 until the day event k occurs, then increasing day after day until the next even occurs, and switching back to 0 when the next event has happened.¹⁶ The associated parameters $\pi_{5ki}^p = \pi_{5k}^p + r_{5ki}^p$ indicate the degree to which the event alters the slope π_{1i}^p of the linear effect of time within constituencies after event k , while the parameter π_{6k}^p indicates the extent to which the event alters the quadratic effect of time estimated by π_2^p for all constituencies together.

Next, we estimate the average effect of changes in similarity induced by the last event occurred before the electoral outcome (π_{3ip}^{11}). This average effect is the change in similarity that occurred based on the days from the event to the Federal election, relative to the counterfactual pre-event trend. Hence, it represents whether the difference between predicted similarity to a party after eleven events happened, and the counterfactual similarity in case no event had happened, is correlated with the electoral outcome. We pool all parties together and estimate

$$\Delta vote_{ip} = \alpha + \beta \pi_{3ip}^{11} + v_{ip} \quad (3)$$

where the vector $\Delta vote_{ip}$ contains the difference in vote share for party p in constituency i between the general elections of 2017 and 2013.

Differently from papers that correlate party votes with economic variables such as unemployment, we correlate votes to change in language similarity. Note that, differently from variables such as unemployment that are fixed at the constituency level, our right-hand side variable can vary across parties within a constituency. Hence, whereas it is not possible to use macroeconomic variables as independent variables when pooling all parties together (because the independent variable does not vary within a constituency, while the dependent variable does), we can use π_{3ip}^{11} thanks to its variation within a constituency.

¹³A series of Log-Likelihood Ratio tests indicate that the inclusion of a quadratic effect of Time improves the fit of each party model (90 percent significance level for all parties, although for most parties we find a much higher significance level). Results are presented in Table E.1 in appendix E.

¹⁴We omit the random coefficients of the quadratic term of *Time*, since models including these random coefficients do not converge.

¹⁵We include only one year dummy due to high multicollinearity in our estimation. Recall that the relevant period is about two years long, from July 2015 until September 2017.

¹⁶For an illustration of the coding of the variables see Table D.1 in appendix D.

As mentioned above, although events occur independently of local characteristics, their effect on similarity could depend on local conditions. However, it is not straightforward to imagine constituency characteristics that could cause between-constituency variation in similarity shifts after an event. We investigated whether a set of standard variables often considered in explaining the growth of populist parties (e.g. unemployment, share of employees working in manufacturing or foreign population) can explain the cross-constituencies heterogeneity but did not find any significant effect. Results are presented in Table E.4 in Appendix E.

Who moves: the Parties or the Public?

One natural concern with our empirical strategy comes from the specific measure of language similarity that we use. We can think of similarity as an equilibrium outcome generated by the interaction between two agents: the party account and the public. In interpreting our results we treat the parties' language on social media as exogenous and assume individuals are getting "closer" or "farther" from the language of different parties according to their shifting views. This assumption would be threatened if parties (AfD in particular) changed the language of their tweets to make it more similar to that of German Twitter users, as a political communication strategy. We do not deny that such or similar mechanisms are part of party strategies on social media. What we are concerned about is: do parties themselves significantly change their communication when events happen? If so, what we argue to be a public shift closer to or farther from a party after specific events could be simply due to party language changing on those days.

To shed light on this issue we aggregate all the tweets that a certain party posts in a week. The weekly aggregation is useful for example to avoid discontinuities caused by party-specific daily events, in comparison to a long-term shift in language use. Then, using the same doc2vec model described above, we compute the within-party change in language similarity on a weekly rolling basis. Finally, we estimate the discontinuous growth model introduced above to see whether the within-party similarity changes around events. In case a party used significantly different language from one week to another in weeks after an event, we would observe a sharp downward shift as the similarity in these weeks would be low. If instead following an event the party keeps using very similar language, we would expect an upward shift in the observed language similarity at the time a party focuses on one specific topic as the similarity will be quite high. Whether in one case or the other, discontinuities in within-party similarity would appear.

Results

Shifts in Similarity

We start presenting our results in Figure 4 and Figure 5. For $k = 1, \dots, 11$ and different parties p we show the estimated coefficients $\pi_{4,k}^p$ (fixed component, uniform across constituencies) representing the difference between the predicted levels in absence of events and the predicted values produced by our model which incorporates discontinuities. Estimates for all parameters of the discontinuous growth model can be found in Table E.2 in Appendix E.

The parties shown in Figure 4 are the AfD and, for comparison, the traditional center-left SPD. Figure 4a shows that changes in language similarity at events is positive and significant for the AfD, while negative and significant for the SPD. Notice that the increasing magnitude of the shifts in similarity of both parties does not stem from an accumulation of coefficient, but from the functional form of the discontinuous growth model which estimates the changes relative to the trend which existed before the first event occurred.

Figures 4b and Figure 4c show, for the two parties, whether there is evidence for the public moving closer to the AfD and farther away from the SPD. As discussed above, a different interpretation could be that the parties change their language in response to events, while the public does not move in either direction. In Figures 4b and Figure 4c observe that, in response to events, the AfD does not change its language, while the SPD does. Combining these observations with finding in Figure 4a, under the relatively weak assumption that the left-wing SPD did not adapt a right-wing language following these events, we can conclude that the public shifted towards AfD as a response to the events. This might partially help explaining the large loss for the SPD at the 2017 general election.

Consider other parties in Figure 5, the results reveal an interesting, and partially unexpected, pattern. AfD is the party that gains the most as we observe increasingly positive similarity shifts at each event. The CSU, the Bavarian ally of Angela Merkel's CDU, traditionally the most right-wing party before the emergence of the AfD, also shows positive shifts in language similarity, although much lower compared to the AfD. This is consistent with the recent party history: the union of CDU and CSU was under enormous pressure during the peak of the refugee crisis around 2015. High-ranked CSU officials challenged Angela Merkel's leadership after she announced an open-border policy for asylum seekers, and started promoting closed borders and deportation.¹⁷ Thus, observing a positive shift in language similarity for this party and the AfD is not surprising. We find positive similarity shifts (albeit marginally significant) also in the case of the two left parties, The Left or The Greens. We observe instead negative similarity shifts at events just for centrist

¹⁷See Foreign Policy (22.10.2015).

parties CDU, FDP, and SPD. Collectively, these results suggest that terror events can increase polarization. Some people move closer (at least in their language) to right-wing parties, mostly AfD; others move closer to left parties The Left or The Greens even though this shift is smaller than the shift to the right. And in general people appear to move farther away from center parties such as CDU, FDP and SPD.

Shifts in Similarity and Votes

We found in the previous section that the events we considered can affect the language distance from right-wing parties relative to center and left parties. We now investigate whether change in similarity can predict electoral outcomes in the 2017 federal election.


Results are presented in Table 4. The dependent variable is the change in vote share from 2013 to 2017 across parties and constituencies. The independent variable is the vector of shifts in similarity to parties across parties and constituencies: $\{\pi_{4,11}\}_{ip}$ for all constituencies i and all parties p . Remember that these shifts are constituency-specific in that we allowed for random coefficients (see Equation 2) As mentioned before, all events are exogenous to local conditions, which are usually measured with standard macroeconomic variables. We are not trying to assess which local characteristics explain electoral outcomes: what we want to investigate is whether cultural incidents and terror events have independent explanatory power for electoral outcomes, beyond other factors orthogonal to those events.

We observe a positive and significant relationship between shifts in similarity and electoral outcomes for all parties combined. In this case, we run a single regression for all parties rather than one for each party because the sample size would be too small. After the large differences presented in Figure 4a and in Figure 5, where AfD appears to be the party with the strongest upward shift, this should not be a surprise considering that AfD was the party with the largest increase in vote share. Estimating this model party by party, however, we do not find a significant correlation.¹⁸ Note however that, as already mentioned, the party by party regression suffers from a low sample size. Although the relationship of Table 4 suffers from the aggregation of votes across parties we emphasize that the coefficient is highly significant.


Although our events are exogenous to any local characteristic, one would still like to know which local characteristics amplify or dampen similarity shifts at the time of events. As explained before, identifying the right set of independent variables that could possibly be correlated with this effect is not obvious. We have considered the set of structural variables identified by Franz et al. (2018) to be appropriate to explain the rise of the AfD in Germany.¹⁹ Note that this set might not

¹⁸Estimation results are available from the authors on request.


¹⁹According to these authors the chosen variables, such as the share of craftsmen firms and other demographic and standard economic variables, are sufficient to represent the socio-economic and demographic conditions in a typical German electoral constituencies.



(a) Shifts in Similarity: AfD and SPD




(b) Within-party Shifts in Similarity: AfD




(c) Within-party Shifts in Similarity: SPD

Figure 4: Shifts in Similarity: AfD vs SPD


Notes: Subfigure 4a shows estimated coefficients $\pi_{4,k}^P$ (fixed component, see Equation 1) for parties AfD and SPD. Confidence interval corresponds to the 95 percent significance level. Subfigures 4b and 4c show point estimates of event specific shifts in intercept, similar to $\pi_{4,k,t}^P$ in Equation 2, as part of the within-party discontinuous growth model estimated for AfD and SPD. Within-party similarity is calculated on a rolling weekly basis. Drawn confidence intervals are for the 90 percent significance level.




(a) CDU




(b) CSU



(c) The Greens



(d) The Left



(e) FDP

Figure 5: Shifts in Similarity: other Parties

Notes: Estimated coefficients $\pi_{4,k}^p$ (fixed component, see Equation 1) for parties CDU, CSU, The Greens, The Left, and FDP. Confidence interval corresponds to the 95 percent significance level.

Table 4: Electoral Effect: Votes on Shifts in Similarity

	Δ Vote Share
Shifts in Similarity	0.0054*** (0.0001)
Constant	0.0033*** (0.001)
N	1350
R^2	0.232

Notes: Δ Vote Share refers to the difference in electoral results between 2017 and 2013. All standard errors are clustered on constituency level and calculated using bootstrapping. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

be the optimal set for all parties. In any case, we do not find any of the considered variables to be correlated with the size of reaction of events. Results are reported in Table E.4 in Appendix E.

Conclusions

The rise of far-right populist parties is at the core of political and scholarly debate in Western democracies. In the European context, these parties are nationalist, protectionist, and often culturally conservative. The observed combination, in their manifestos, of opposition to free circulation of goods and people, along with rejection of multiculturalism and security concerns related to immigration, has led political scientists and economists to try and disentangle economic and cultural factors behind the populist vote.

In this paper we have exploited the exogenous timing of salient events of terrorism and crime to look for an effect on peoples' alignment with the values promoted by a right-wing populist party. Using an allocation rule based on following patterns of local accounts to assign Twitter users to geographic constituencies, and a deep learning model, we show that unexpected terrorist attacks and a salient crime event move the language of peoples' tweets closer to that of the German right-wing party.

Our interpretation is that terrorist attacks and large-scale crimes attributed to immigrants constitute shocks that dramatically increase the salience of cultural differences and elicit perceptions of threats and hostility from a different religion. To some degree, these concerns might have always been present in the population but are sparked by the events we study, possibly moving the political leaning of Twitter users towards the party which emphasized threats from immigration and multiculturalism the most.

Our evidence suggests that terrorist attacks and other cultural incidents correlate with votes. We find that following the occurrence of major terrorist or culturally charged events, the average electoral constituency observes a significant and lasting shift in the language of its tweets towards the language of far-right AfD. The same constituencies shift away from the centrist left SPD and

the centrist CDU, the two main parties at the time in government. Moreover, we find some weak but suggestive evidence that similarity increases for the far-left parties, indicating that terrorist attacks may increase polarization following a sequence of events.

One should be cautious in interpreting our results as causal. Nonetheless, our evidence suggests that party behaviors do not drive our findings, which are consistent with the channel we hypothesize.

Overall, these findings advance our understanding of the roots of radical right support, stressing the role of perceived cultural threats elicited by terrorist events and culturally salient crimes. They also contribute to the literature on the effects of terrorism on public opinion and elections, by showing that attacks have an effect on the support for parties promoting isolationism and cultural conservatism. Moreover, they highlight a significant connection between online behavior and political outcomes, confirming the relevance of social media text as a measure of attitudes. They also show the potential of using information from individual accounts' following patterns to locate geographically social media users and exploit cross-sectional variation in their distribution in empirical designs.

Our case-study is limited to a single country. However, given the concurrent surge of radical right and terrorism in several Western democracies, we believe these results are relevant to other geographical settings. Moreover, the combination of exogenous real world events and geo-referenced social media data is a promising approach for other areas of social science. For instance, it might be possible to study how people react online and offline, in the short and medium term, to crime events happening in their proximity (Mobasseri 2019). Another possibility could be to bring the study of online behavior in the aftermath of terrorist events to areas where different ethnic or national groups co-exist and relate it to integration or discrimination outcomes. Exploring these ideas further is an exciting avenue for future research.

References

- Baccini, Leonardo and Weymouth, Stephen (2019), Gone For Good: Deindustrialization, White Voter Backlash, and U.S. Presidential Voting. *Working paper*.
- Ballard-Rosa, Cameron, Jensen, Amalie, and Scheve, Kenneth (2018), Economic decline, social identity, and authoritarian values in the United States. *American Political Science Association, Boston*.
- Barberá, Pablo (2015), Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. *Political Analysis*. 23(1), 76–91.

- Barberá, Pablo, Casas, Andreu, Nagler, Jonathan, Egan, Patrick J, Bonneau, Richard, Jost, John T, and Tucker, Joshua A (2019), Who leads? Who follows? Measuring issue attention and agenda setting by legislators and the mass public using social media data. *American Political Science Review*. 113(4), 883–901.
- Barone, Guglielmo, D’Ignazio, Alessio, Blasio, Guido de, and Naticchioni, Paolo (2016), Mr. Rossi, Mr. Hu and politics. The role of immigration in shaping natives’ voting behavior. *Journal of Public Economics*. 136, 1–13.
- Berrebi, Claude and Klor, Esteban F (2008), Are voters sensitive to terrorism? Direct evidence from the Israeli electorate. *American Political Science Review*. 102(3), 279–301.
- Bhatia, Shraey, Lau, Jey Han, and Baldwin, Timothy (2016), Automatic Labelling of Topics with Neural Embeddings. *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 953–963.
- Bisbee, James (2019), What Is Out Your Back Door: How Policy Preferences Respond to Local Trade Shocks. *Working paper*.
- Bliese, Paul D. and Lang, Jonas W. B. (2016), Understanding Relative and Absolute Change in Discontinuous Growth Models. *Organizational Research Methods*. 19(4), 562–592.
- Cantoni, Davide, Hagemester, Felix, and Westcott, Mark (2019), Persistence and Activation of Right-Wing Political Ideology. *Working Paper*.
- Cerrato, Andrea, Ferrara, Federico Maria, and Ruggieri, Francesco (2018), Why Does Import Competition Favor Republicans? *Working paper*.
- Colantone, Italo and Stanig, Piero (2018a), Global Competition and Brexit. *American Political Science Review*. 112(02), 201–218.
- (2018b), The Economic Determinants of the ‘Cultural Backlash’: Globalization and Attitudes in Western Europe. *Working Paper*. (2018-91).
- (2018c), The trade origins of economic nationalism: Import competition and voting behavior in Western Europe. *American Journal of Political Science*. 62(4), 936–953.
- Corriere Della Sera (28.03.2018), Salvini: Stop a tutti gli sbarchi, il rischio terrorismo è altissimo https://www.corriere.it/politica/18_marzo_28/salvini-stop-tutti-sbarchi-rischio-terrorismo-altissimo-c3ccda52-3280-11e8-99d9-35360f097e9f.shtml.
- Couttenier, Mathieu, Hatte, Sophie, Thoening, Mathias, and Vlachos, Stephanos (2019), The Logic of Fear-Populism and Media Coverage of Immigrant Crimes. *Working Paper*.
- Curini, Luigi, Iacus, Stefano, and Canova, Luciano (2015), Measuring Idiosyncratic Happiness Through the Analysis of Twitter: An Application to the Italian Case. *Social Indicators Research*. 121(2), 525–542.

- Darius, Philipp and Stephany, Fabian (2019), "Hashjacking" the Debate: Polarisation Strategies of Germany's Political Far-Right on Twitter. *International Conference on Social Informatics*. Springer, pp. 298–308.
- De Smedt, Tom and Jaki, Sylvia (2018), The Polly Corpus: online political debate in Germany. *Proceedings of CMC and Social Media Corpora*, 33–36.
- Der Bundeswahlleiter (2017), *Wahlkreise und zugeordnete Gemeinden bei der Wahl zum 19. Deutschen Bundestag*. Wiesbaden.
- Dustmann, Christian, Vasiljeva, Kristine, and Damm, Anna Piil (forthcoming), Refugee migration and electoral outcomes. *Review of Economic Studies*.
- Eady, Gregory, Nagler, Jonathan, Guess, Andy, Zilinsky, Jan, and Tucker, Joshua A (2019), How Many People Live in Political Bubbles on Social Media? Evidence From Linked Survey and Twitter Data. *SAGE Open*. 9(1), 2158244019832705.
- Echebarria-Echabe, Agustin and Fernández-Guede, Emilia (2006), Effects of terrorism on attitudes and ideological orientation. *European Journal of Social Psychology*. 36(2), 259–265.
- Eurostat (2019), *Asylum Statistics*. https://ec.europa.eu/eurostat/statistics-explained/index.php/Asylum_statistics.
- Ferrara, Emilio, Varol, Onur, Davis, Clayton, Menczer, Filippo, and Flammini, Alessandro (2016), The Rise of Social Bots. *Communications of the ACM*. 59(7), 96–104.
- Ferrín, Mónica, Mancosu, Moreno, and Cappiali, Teresa M. (2019), Terrorist attacks and Europeans' attitudes towards immigrants: An experimental approach. *European Journal of Political Research*.
- Finseraas, Henning, Jakobsson, Niklas, and Kotsadam, Andreas (2011), Did the Murder of Theo van Gogh Change Europeans' Immigration Policy Preferences? *Kyklos*. 64(3), 396–409.
- Foreign Policy (22.10.2015), Nothing Can Take Down Angela Merkel — Except 800,000 Refugees.
- Franz, Christian, Fratzscher, Marcel, and Kritikos, Alexander S. (2018), German right-wing party AfD finds more support in rural areas with aging populations. *DIW Weekly Report, DIW Berlin*.
- Frees, Beate and Koch, Wolfgang (2018), ARD/ZDF-Onlinestudie 2018: Zuwachs bei medialer Internetnutzung und Kommunikation: Ergebnisse aus der Studienreihe "Medien und ihr Publikum" (MiP). *Media Perspektiven*. (9), 398–413.
- Gentzkow, Matthew and Shapiro, Jesse M. (2010), What Drives Media Slant? Evidence From U.S. Daily Newspapers. *Econometrica*. 78(1), 35–71.
- Getmansky, Anna and Zeitzoff, Thomas (2014), Terrorism and voting: The effect of rocket threat on voting in Israeli elections. *American Political Science Review*. 108(3), 588–604.
- Goldstein, Harvey (2013), *Multilevel Statistical Models (3rd ed.* London: Arnold.

- Gould, Eric D and Klor, Esteban F (2016), The Long-run Effect of 9/11: Terrorism, Backlash, and the Assimilation of Muslim Immigrants in the West. *The Economic Journal*. 126(597), 2064–2114.
- Halla, Martin, Wagner, Alexander F., and Zweimüller, Josef (2017), Immigration and voting for the far right. *Journal of the European Economic Association*. 15(6), 1341–1385.
- Hangartner, Dominik, Dinas, Elias, Marbach, Moritz, Matakos, Konstantinos, and Xefteris, Dimitrios (2017), Does Exposure to the Refugee Crisis Make Natives More Hostile? *American Political Science Review*, 1–14.
- Hansen, Michael A. and Olsen, Jonathan (2019), Flesh of the Same Flesh: A Study of Voters for the Alternative for Germany (AfD) in the 2017 Federal Election. *German Politics*. (28), 1–19.
- Hill, Seth J, Hopkins, Daniel J, and Huber, Gregory A (2019), Local demographic changes and US presidential voting, 2012 to 2016. *Proceedings of the National Academy of Sciences*. 116(50), 25023–25028.
- Infratest Dimap (2018), *Umfragen & Analysen: Bundesländer*.
- Inglehart, Ronald and Norris, Pippa (2016), Trump, Brexit, and the Rise of Populism: Economic Have-Nots and Cultural Backlash. *Working Paper*.
- Jardina, Ashley (2019), *White identity politics*. Cambridge University Press.
- Jones, Kevin L, Noorbaloochi, Sharareh, Jost, John T, Bonneau, Richard, Nagler, Jonathan, and Tucker, Joshua A (2018), Liberal and conservative values: What we can learn from congressional tweets. *Political Psychology*. 39(2), 423–443.
- Jost, John T, Barberá, Pablo, Bonneau, Richard, Langer, Melanie, Metzger, Megan, Nagler, Jonathan, Sterling, Joanna, and Tucker, Joshua A (2018), How social media facilitates political protest: Information, motivation, and social networks. *Political psychology*. 39, 85–118.
- Kibris, Arzu (2011), Funerals and elections: The effects of terrorism on voting behavior in Turkey. *Journal of Conflict Resolution*. 55(2), 220–247.
- Kobold, Kevin and Schmiedel, Sven (2018), Wahlverhalten bei der Bundestagswahl 2017 nach Geschlecht und Alter: Ergebnisse der repräsentativen Wahlstatistik. *WISTA - Wirtschaft und Statistik*. (3), 142–156.
- Larson, Jennifer M., Nagler, Jonathan, Ronen, Jonathan, and Tucker, Joshua A. (2019), Social networks and protest participation: Evidence from 130 million Twitter users. *American Journal of Political Science*. 63(3), 690–705.
- Le, Quoc and Mikolov, Tomas (2014), Distributed representations of sentences and documents. *International conference on machine learning*, pp. 1188–1196.
- Legewie, Joscha (2013), Terrorist Events and Attitudes toward Immigrants: A Natural Experiment. *American Journal of Sociology*. 118(5), 1199–1245.

- McConnell, Brendon and Rasul, Imran (2018), Contagious animosity in the field: Evidence from the Federal Criminal Justice System. *Working Paper*.
- Mikolov, Tomas, Chen, Kai, Corrado, Greg, and Dean, Jeffrey (2013), Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mobasseri, Sanaz (2019), Race, place, and crime: how violent crime events affect employment discrimination. *American journal of sociology*. 125(1), 63–104.
- Müller, Karsten and Schwarz, Carlo (2019), Fanning the flames of hate: Social media and hate crime. *Working paper*.
- (2020), From hashtag to hate crime: Twitter and anti-minority sentiment. *Working paper*.
- Mutz, Diana C. (2018), Status threat, not economic hardship, explains the 2016 presidential vote. *Proceedings of the National Academy of Sciences of the United States of America*. 115(19), E4330–E4339.
- Schüller, Simone (2016), The effects of 9/11 on attitudes toward immigration and the moderating role of education. *Kyklos*. 69(4), 604–632.
- Singer, Judith D. (1998), Using SAS PROC MIXED to Fit Multilevel Models, Hierarchical Models, and Individual Growth Models. *Journal of Educational and Behavioral Statistics*. 23(4), 323–355.
- Steinmayr, Andreas (2018), Contact Matters: Exposure to Refugees and Voting for the Far-Right. *Working paper*.
- Tabellini, Marco (forthcoming), Gifts of the immigrants, woes of the natives: Lessons from the age of mass migration. *Review of Economic Studies*.
- Vaccari, Cristian, Valeriani, Augusto, Barberá, Pablo, Bonneau, Rich, Jost, John T, Nagler, Jonathan, and Tucker, Joshua A (2015), Political expression and action on social media: Exploring the relationship between lower-and higher-threshold political activities among Twitter users in Italy. *Journal of Computer-Mediated Communication*. 20(2), 221–239.
- Vaccari, Cristian, Valeriani, Augusto, Barberá, Pablo, Bonneau, Richard, Jost, John T., Nagler, Jonathan, and Tucker, Joshua (2013), Social media and political communication. A survey of Twitter users during the 2013 Italian general election. *Rivista italiana di scienza politica*.
- Zeit Online (5.01.2016), Polizei gibt Fehler in Silvesternacht zu.
<https://www.zeit.de/gesellschaft/zeitgeschehen/2016-01/uebergriffe-koeln-frauen-silvester-henriette-reker>.
- (24.02.2017), "Wir können uns nicht von Kinderaugen erpressen lassen"
<https://www.zeit.de/politik/deutschland/2016-02/alexander-gauland-afd-fluechtlingskrise-fluechtlingspolitik-grenzen>.

Appendices

A Appendix: Tweets and Content	31
B Appendix: Text Processing Details	33
C Appendix: Validation	35
D Appendix: Discontinuous Growth Model	38
E Appendix: Results	39

A Appendix: Tweets and Content

Parties' Tweets

In the following we first provide the details on how wordclouds are created, and then present the visualization between the two parties The Greens and AfD. To construct wordclouds we first compute the following log-odds-ratio for each word²⁰ w in the tweets of a party:

$$\log or_w = \log \left[\frac{f_{i,w}}{1 - f_{i,w}} \right] - \log \left[\frac{f_{j,w}}{1 - f_{j,w}} \right]$$

where $f_{i,w}$ is the frequency of a word w in document i .²¹ This ratio identifies words which are most likely to appear in a party's tweets and at the same time least likely in another party's tweets, thus allowing us to identify what one party is most concerned about but the other one is not. We then rank them from highest to lowest and categorize the resulting list. Some of these words naturally occur mainly in one of the parties' tweets but not in another: for instance the names of politicians and party specific events, such as congresses and internal elections. Since we are mainly interested in identifying the words with political relevance we manually categorize each word, such that we know whether it is about a political topic or about something else, like the name of a politician, the reference to an event or non-identifiable junk.

After obtaining a categorized list ordered by the ratio shown above, we focus on the words with a political meaning with a minimum of three occurrences and create a "wordcloud" out of them.²² Here, the size of a word reflects its frequency of occurrence in the respective document among those words which are at the same time least likely to occur in the other document.

We already compared in Figure 2 in the main text two wordclouds for the AfD, computed before and after July 2015. Here, Figure A.1 shows a comparison between the language of the right-wing AfD and the one of the left party The Greens for all years in which we collect their tweets. The picture for the AfD shows a mix of economic topics, mainly concerned with the European debt crisis (*ecb*, *euro_crisis*, *greek*, *monetary_policy*). On the other hand, the picture for The Greens shows the expected focus on environmental topics (*factory_farming*, *plastic_waste*, *climate_goal*) mixed with traditional leftist topics (*ttip_protest*, *refugees_welcome*, *female_quota*).

Constituency Tweets: Racist Words and Votes

Besides analyzing the language of parties, we also analyzed the language used in constituencies. Recall that we use the corpus of tweets containing hate speech provided by De Smedt et al. (2018)

²⁰With the term "word", we actually mean a "token" after pre-processing the tweets, as explained below in appendix section B.

²¹We use the log normalization to make the odds-ratios symmetric across documents.

²²All phrases and words in Figure A.1 and Figure 2 are translated from German into English.

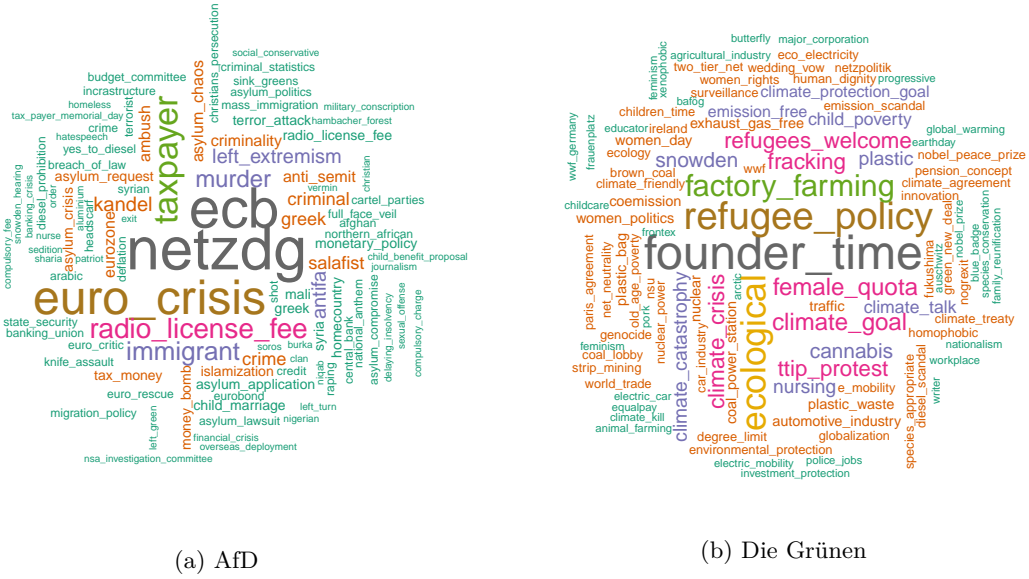


Figure A.1: Between Party Comparison: AfD vs The Greens

to manually extract a list of right-wing leaning or racist words, for which we then turn to the constituencies to calculate the share of usage of these words (out of the total amount of words used). Table A.1 presents the bivariate regression results using the difference in electoral results between 2017 and 2013 for each party as a dependent variable and the share of racist words (out of the total number of words) in a constituency as independent variable. We find that the share of racist words is positively and significantly correlated with the change in votes for AfD. Furthermore, we find no significant correlation with other parties except for the most left party (The Left), which shows a negative significant correlation.

Table A.1: Racist Words and Votes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	AfD	CDU	CSU	The Greens	FDP	The Left	SPD
Racist Words (Share)	5.75* (2.48)	-1.59 (1.61)	6.93 (14.79)	-0.48 (0.67)	-0.12 (0.79)	-3.32* (1.54)	-1.35 (1.01)
Constant	0.08*** (0.00)	-0.08*** (0.00)	-0.11*** (0.01)	0.00*** (0.00)	0.06*** (0.00)	0.01*** (0.00)	-0.05*** (0.00)
Observations	232	195	37	232	232	232	232

Notes: Dependent variable the difference in electoral results between 2017 and 2013. Racist words extracted from the corpus of tweets provided by De Smedt et al. (2018). See also related discussion in the main text. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B Appendix: Text Processing Details

As discussed in the main text, we compute similarity between the tweets of the parties and the tweets of each constituency by transforming the two groups of tweets into vectors using `doc2vec`, a deep learning model that we describe below. We then measure similarity as the cosine similarity between the two vectors. Before proceeding with `doc2vec`, we pre-process tweets. In the following paragraphs we provide details on these steps.

Text Preprocessing

Text pre-processing is necessary to reduce the computational time necessary to run the `doc2vec` model. Computational time is more than directly proportional to vocabulary size, namely the number of words in our corpus of tweets. With pre-processing we reduce the number of words, and hence computational time, without losing relevant information. We follow standard procedures in text pre-processing with different libraries in Python. First we lower-case all words and tokenize the text, i.e., we break streams of text into single words, called “tokens”. We do this using “`word_tokenize`” from the Python module `NLTK`. Next, we eliminate punctuation and stopwords, namely words that recur very frequently in our corpus and have little meaning. The dictionary of stopwords we use is the one in `NLTK`. We also remove all tokens that consist of non-alphanumeric characters only, and remove emoticons, links, `@`, and `#` symbols. Then, we perform “stemming”, which implies conflating the variant forms of a word into a common representation, the stem. For instance, the words “ate” and “eating” are both reduced to the common stem “eat”. Stemming relies on existing dictionaries: we use the German Stemmer in the Python module “`gensim`”. Finally, we perform collocations, namely, we identify combinations of two words that have a higher probability of occurring together than separately. For instance, the tokens “angela” and “merkel” have higher chances of co-occurring as the bigram “angela merkel” than separately. In this case, collocations transform the two separate tokens into just one: “angela_merkel”. We used `Bigram CollocationFinder` in `NLTK`. We then use the pre-processed tweets to train the `doc2vec` model.

`doc2vec`

After pre-processing our tweets, we create two “documents”, each at daily frequency: a party- and a constituency- document. The party document is the text of all the tweets the party posted on a certain day. A constituency document is the text of all the tweets that all the users in our sample located in a given constituency posted on a certain day. Since we have 752 days in our observation period (from September 4th 2015 to September 24th, 2017,²³ we end up with 752 documents for

²³As stated in the main text, July 2015 marked a turning point in the history of the AfD. We leave two months between the change in leadership of the AfD and the starting point of our analysis, but we emphasize that the

each party and 752 documents for each constituency in our sample.²⁴

We use doc2vec (Le et al. 2014), an unsupervised deep learning algorithm that learns how to represent each document with a unique vector, and which is a generalization of Word2Vec. In order to understand doc2vec it is necessary to first understand how Word2Vec works. Word2Vec (Mikolov et al. 2013) is an unsupervised deep learning algorithm that learns how to represent each word as a vector, depending on the surrounding (context) words. It takes as input a large vocabulary of words, trains a neural network language model with a single hidden layer, and produces a vector space, where each word is represented as a vector in this space. Word vectors, also called word embeddings, are positioned in the vector space such that words with similar semantic meaning are located in close proximity to one another. The model is trained using stochastic gradient descent with back propagation. When the algorithm converges, it represents words as word embeddings, namely meaningful real-valued vectors of configurable dimension (usually, 300 dimensions).

doc2vec (Le et al. 2014) is an extension of Word2Vec which learns to represent not just individual words, but entire documents. By treating each document as a word token, the same Word2Vec methodology is used to learn document embeddings (Bhatia et al. 2016). As in Word2Vec, training happens through back propagation through several iterations. Each iteration of the algorithm is called an “epoch”, and its purpose is to increase the quality of the output vectors. This type of document embedding allows to represent texts as dense fixed-length feature vectors, taking into account their semantic and syntactic structure.

We use the Distributed Bag of Words (DBOW) model and a freely available implementation of the doc2vec algorithm included in the gensim Python module, whose implementation requires the following hyperparameters:

- Size: the dimensionality of the vector representing the document. We set it to 300.
- Window size: The maximum distance between the current and predicted word within a sentence. We set it to 15.
- Epochs: Number of iterations over the corpus to train the algorithm. We set it to 300.
- Min_count: Ignores all words with total frequency lower than this. We set it to 20.
- Sub-sampling: The threshold for configuring which higher-frequency words are randomly down-sampled: useful range is $(0, 10^{-5})$. We set it to 10^{-3} .
- Negative: The number of “noise words” that should be drawn. We set it to 5.

With the resulting measures, we compute the cosine similarity described in the main text.

empirical method chosen is not sensitive to the exact day. More importantly, we choose this day as it represents a major peak of events during the 2015 refugee crisis with the German government opening borders for refugees coming from Hungary.

²⁴752 is the maximum possible amount for a constituency whose users posted tweets every single day.

C Appendix: Validation

In order to control whether the computed similarity is indeed a valid measurement for how close public opinion is to the various parties, in this section we perform three validation experiments. These experiments first perform a basic but intuitive check by comparing the performance of doc2vec in assessing how close two documents of texts are, with the assessment by a human reader. Next, we focus on parties and assess whether the results produced by the doc2vec algorithm match common priors about the use of language by different parties. We use the corpus of German tweets discussed above containing racist rhetoric and find that the AfD is indeed the party using language closest to it. Finally, we focus on the German public and use public opinion and electoral data and observe that they also are correlated with our measure of similarity.

Human Reader

Tweets of German parties and constituencies from four different days were assessed by a native German who first read both the tweets within a constituency and the corresponding party tweet on a given day and then judged between high or low similarity. The answers confirm the high and low text similarity computed by the doc2vec algorithm.²⁵ This validation is only on a basic level, and is no proof of a valid similarity, but it presents a first transparent way to assess the quality of our measurement.

Racist Tweets

The corpus of German tweets by far-right supporters discussed in the main text offers another possibility to validate the doc2vec performance. Given that many tweets inside the corpus contain violent (and possibly illegal) rhetoric, often with clear references to the Third Reich, we would not expect a very close alignment with any of the official national Twitter accounts of German parties, independent of their ideology. Nevertheless, given that many of the tweets contained in this corpus also use what one can consider current right-wing rhetoric, we would expect at least a closer alignment of the tweets with the language of a right-wing party than with the one of a left-wing party.²⁶ Observing the opposite would raise serious doubts on the effectiveness of doc2vec when applied to public opinion.

We calculate the weekly similarity between the language used by an official party account and the corpus of right-wing tweets. Here we move from a daily to weekly level for two reasons. First,

²⁵Tweets and similarities available from the authors on request.

²⁶In order to emphasize this point, consider the following examples. While *Islamisierung* (islamization), *norefugees*, and *Überfremdung* (foreign domination) are common phrases of right-wing party rhetoric and can also be found many times within the corpus of racist tweets, more extreme language, such as *NS jetzt* (national socialism now), *Waffen SS* or *Menschenmüll* (human trash), cannot be found in the official language used on Twitter by any main party account.

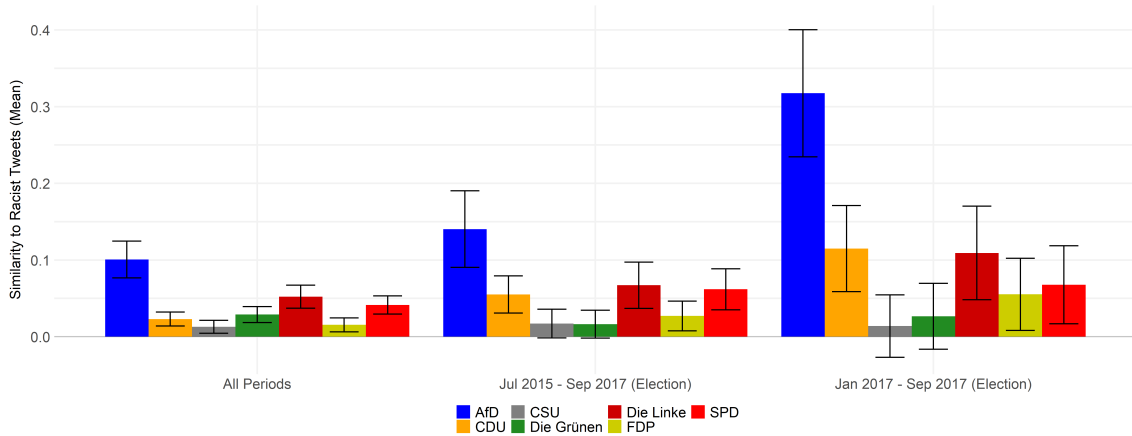


Figure C.1: Similarity to Racist Tweets

Notes: Bars show the mean similarity between the corpus of racist tweets and the language of different parties for the specified time periods. Similarity is calculated at the weekly level. Intervals represent the 95 percent confidence interval of a t-test for the null hypothesis of a respective mean being equal to zero.

as this is not an event study, we are less restricted to calculate the similarity on specific days. Second and more importantly, since we are not estimating an econometric model which requires a considerable amount of data even for subsets of periods, the reduction of the sample size when switching from daily to weekly is not as problematic as it would be at a later stage when we require as many observations between events as possible. What we gain by moving to the weekly level are units representing larger bodies of text, compared to the daily level, which is helpful when focusing only on single accounts.

After obtaining the similarity between the parties and racist tweets, we perform a mean comparison for different sub-periods. The result, together with the 95 percent confidence interval of a t-test for the null hypothesis of a mean being equal to zero, is visualized in Figure C.1. Two observations are noteworthy. First, the mean similarity of the AfD to racist tweets is the highest in all of the sub-periods considered. At first glance the mean itself might not appear to be very high. However, as we said, this corpus of racist tweets partially includes much more extreme language than can be attributed to any official party account, including the AfD. Second, the mean of several parties seems to be increasing if we narrow the time window towards the 2017 election day, probably in an effort to mobilize the supporters closer to the election. But, as expected, this increase is largest for the AfD. In conclusions, this analysis seems to support the validity of our doc2vec model, whose results are consistent with common expectations.

Votes and Polls

To further control the validity of the results of our doc2vec model, we compute the correlation between our measure of similarity and a) the results of the 2017 federal election at constituency

level, and b) poll data provided by Infratest Dimap (2018) at state level.

For the election outcomes, we merge the tweets posted in the 30 days before the election within electoral constituencies and then apply the doc2vec algorithm. We repeat this 15 days before the election as a robustness check. The reason for merging texts over 30 days is to produce a sufficient amount of text for both parties and constituencies as not all parties posted tweets in the days immediately before the election. For the analysis of poll data, since poll surveys are conducted at state level, we merge the tweets of all the constituencies in a given state on the day of the poll.

We then perform two regression analyses: one with the change in vote share from 2013 to 2017 for all parties as the dependent variable, and one with the poll results as the dependent variable.²⁷ In both cases, we regress the dependent variables on the measured similarity. We cluster standard errors on the lowest aggregate for the units of observation, i.e. electoral constituency level or state level, respectively. For the regression on poll results in levels we include party fixed effects to control for variations in levels of party support. Results are presented in Table ???. We observe a positive correlation in all analyses. This analysis offers further support for the fact that our computed similarity captures the public mood across states and electoral borders.

²⁷Each observation is party-constituency (in case change in vote share is the electoral outcome as dependent variable) or party-state-date (in case the poll results at state level on a given day is the dependent variable).

D Appendix: Discontinuous Growth Model

Variables considered include:

1. *Time*: The first variable represents the linear time trend found in a typical growth model.
2. *Time*²: Similar to before with a quadratic time trend.
3. *E*: Event specific change in intercept variable coded 0 prior to the event and 1 after the event, until the next event occurs.
4. *Reset*: Event specific change in slope variable coded 0 at the period of which the event first occurs and increases with each subsequent period until the next event.
5. *Reset*²: Similar to before with a quadratic change variable.

For analyzing multiple events, we simply introduce multiple variables for events and changes. The following table offers an overview on the coding of variables:

Table D.1: Coding of Time Variables - Multiple Events

<i>Time</i>	<i>Time</i> ²	<i>E</i> ₁	<i>E</i> ₂	<i>Reset</i> ₁	<i>Reset</i> ₁ ²	<i>Reset</i> ₂	<i>Reset</i> ₂ ²
0	0	0	0	0	0	0	0
1	1	0	0	0	0	0	0
2	4	0	0	0	0	0	0
3	9	0	0	0	0	0	0
4	16	0	0	0	0	0	0
5	25	1	0	0	0	0	0
6	36	1	0	1	1	0	0
7	49	1	0	2	4	0	0
8	64	1	0	3	9	0	0
9	81	1	0	4	16	0	0
10	100	0	1	0	0	0	0
11	121	0	1	0	0	1	1
12	144	0	1	0	0	2	4
13	169	0	1	0	0	3	9
14	196	0	1	0	0	4	16

Notes: the first event occurs in period 5, the second event in period 10.

E Appendix: Results

The following tables complement the choice of the functional form of the discontinuous growth model (DGM) introduced in equation 1 in the main text, as well as results visualized and discussed in the main text.

Table E.1 shows the results of likelihood ratio tests for the null hypothesis of having a better fit by only including a linear time trend into the discontinuous growth models for each party, compared to the alternative of adding also a quadratic term (no event variables included). As one can see, the null hypothesis is rejected in all cases at the 90 percent significance level, and for most cases even at the 99 percent significance level. We thus include the quadratic time trend in the DGM for all parties.

Table E.2 shows the full list of estimated coefficients for the DGM for each party.

Table E.3 shows the estimated coefficients for the within-party discontinuous growth models for AfD and SPD, visualized in Figure 4b and Figure 4c in the main text.

Finally, Table E.4 shows an investigation at the potential determinants of the heterogeneity in estimated random coefficients of the main DGM. The explanatory variables used corresponds to a set used by Franz et al. (2018) to explain the electoral success of the AfD after the 2017 general election. Notice that the optimal set of explanatory variables may vary across parties, but for the sake of comparison we used the same explanatory variables in each regression.²⁸ As discussed in the main text, we do not find any significant relationship of these variables with the magnitude of the estimated random coefficients.

Table E.1: Likelihood Ratio Test Results: Linear vs Quadratic Time Trend

Party	Likelihood Ratio Test Statistic	p-Value
AfD	804.49	<0.01
CDU	498.4	<0.01
CSU	3.53	0.06
FDP	35.31	<0.01
SPD	56.46	<0.01
The Greens	238.91	<0.01
The Left	3.8	0.05

Notes: Test results refer to a likelihood ratio test for the null hypothesis of a better fit using only a linear time trend in equation 1 versus the alternative of adding a quadratic term.

²⁸Notice that the only variation in the set of explanatory variables is caused by the CSU not existing in the eastern part of Germany, hence the East indicator is excluded.

Table E.2: Discontinuous Growth Model Results

	AfD	CDU	CSU	FDP	SPD	Die Grünen	Die Linke
<i>Time</i>	0.000162 (0.000162)	-0.000167 (0.000202)	-0.000045 (0.000413)	-0.000728 (0.000168)	-0.001486 (0.000172)	0.000635 (0.00018)	-0.000112 (0.000176)
<i>Time</i> ²	-0.000017 (0.000003)	0.000004 (0.000003)	-0.000004 (0.000006)	0.000009 (0.000002)	0.000018 (0.000002)	-0.000005 (0.000003)	-0.000004 (0.000002)
<i>Event</i> ₁	-0.189247 (0.004845)	-0.129345 (0.005022)	0.015701 (0.010167)	0.021558 (0.004344)	0.011961 (0.004335)	-0.020903 (0.004444)	-0.015322 (0.004471)
<i>Event</i> ₂	0.200175 (0.02006)	-0.036653 (0.020331)	0.057938 (0.041784)	-0.071380 (0.016965)	-0.096568 (0.017397)	-0.003214 (0.018158)	0.063682 (0.01783)
<i>Event</i> ₃	0.544628 (0.075778)	-0.112478 (0.076877)	0.159906 (0.157436)	-0.226074 (0.063891)	-0.444399 (0.065531)	0.056927 (0.068506)	0.184817 (0.067287)
<i>Event</i> ₄	1.562022 (0.216935)	-0.341091 (0.220024)	0.378014 (0.450713)	-0.692040 (0.182825)	-1.346701 (0.187531)	0.216244 (0.196067)	0.405169 (0.192563)
<i>Event</i> ₅	3.708360 (0.532385)	-0.848453 (0.539952)	0.878455 (1.106049)	-1.653869 (0.448664)	-3.355275 (0.460201)	0.757909 (0.481154)	0.865480 (0.472556)
<i>Event</i> ₆	5.342222 (0.783078)	-1.287784 (0.794209)	1.309918 (1.626826)	-2.423489 (0.659925)	-4.924571 (0.676899)	1.183823 (0.70772)	1.266131 (0.695069)
<i>Event</i> ₇	5.910504 (0.871105)	-1.458576 (0.883488)	1.418320 (1.809691)	-2.714236 (0.734108)	-5.518408 (0.752989)	1.402249 (0.787277)	1.394901 (0.773203)
<i>Event</i> ₈	6.551063 (0.973689)	-1.573876 (0.987529)	1.607238 (2.022825)	-3.017193 (0.820559)	-6.214220 (0.841664)	1.592895 (0.879989)	1.524770 (0.864259)
<i>Event</i> ₉	6.838220 (1.013608)	-1.648281 (1.028016)	1.637697 (2.105733)	-3.158834 (0.854199)	-6.427453 (0.87617)	1.651038 (0.916066)	1.585600 (0.89969)
<i>Event</i> ₁₀	8.520463 (1.277603)	-2.077992 (1.295765)	2.043185 (2.654166)	-3.993783 (1.076676)	-8.144979 (1.104369)	2.098190 (1.154656)	1.941026 (1.134016)
<i>Event</i> ₁₁	9.271648 (1.393324)	-2.275450 (1.413131)	2.205123 (2.89458)	-4.376816 (1.174198)	-8.913970 (1.2044)	2.279909 (1.259241)	2.085836 (1.236731)
<i>Reset</i> ₁	0.005987 (0.000399)	0.007958 (0.000403)	-0.000962 (0.000822)	-0.002376 (0.00034)	-0.001407 (0.000348)	0.001379 (0.000363)	0.005159 (0.000357)
<i>Reset</i> ₂	0.003423 (0.000501)	-0.001546 (0.000508)	-0.000230 (0.00104)	-0.001468 (0.000423)	-0.002579 (0.000434)	-0.001858 (0.000453)	0.002149 (0.000445)
<i>Reset</i> ₃	0.003900 (0.000929)	-0.001966 (0.000942)	0.000731 (0.001929)	-0.003271 (0.000783)	-0.005860 (0.000803)	0.000059 (0.000839)	0.001799 (0.000824)
<i>Reset</i> ₄	0.010313 (0.001559)	-0.002966 (0.001581)	0.002050 (0.003239)	-0.004997 (0.001314)	-0.009774 (0.001348)	0.003660 (0.001409)	0.002236 (0.001384)
<i>Reset</i> ₅	0.015091 (0.002446)	-0.005741 (0.002481)	0.003973 (0.005082)	-0.007199 (0.002062)	-0.015145 (0.002115)	0.006581 (0.002211)	0.004189 (0.002172)
<i>Reset</i> ₆	0.009788 (0.003046)	-0.001701 (0.00309)	-0.001753 (0.006325)	-0.011162 (0.002569)	-0.022279 (0.002634)	0.013456 (0.002754)	-0.000715 (0.002705)
<i>Reset</i> ₇	0.017480 (0.003181)	0.005051 (0.003227)	0.004213 (0.006608)	-0.010489 (0.002682)	-0.019457 (0.00275)	0.005741 (0.002875)	-0.002691 (0.002824)
<i>Reset</i> ₈	0.020356 (0.004246)	-0.011072 (0.004309)	-0.010721 (0.008804)	-0.023624 (0.003578)	0.010850 (0.00367)	0.003069 (0.003837)	0.001502 (0.003769)
<i>Reset</i> ₉	0.020822 (0.003367)	-0.006735 (0.003415)	0.004702 (0.006994)	-0.010330 (0.002837)	-0.022390 (0.00291)	0.004238 (0.003043)	0.004026 (0.002989)
<i>Reset</i> ₁₀	0.025338 (0.003839)	-0.007092 (0.003894)	0.003990 (0.007981)	-0.013295 (0.003237)	-0.026298 (0.003319)	0.004987 (0.00347)	0.004520 (0.003409)
<i>Reset</i> ₁₁	0.016572 (0.005272)	-0.020155 (0.005353)	-0.001945 (0.010917)	-0.038249 (0.004443)	-0.021849 (0.004557)	0.006977 (0.004765)	0.007718 (0.00468)
<i>Reset</i> ₁ ²	0.000072 (0.000007)	-0.000228 (0.000007)	0.000047 (0.000015)	0.000026 (0.000006)	-0.000024 (0.000006)	-0.000013 (0.000007)	-0.000076 (0.000007)
<i>Reset</i> ₂ ²	0.000015 (0.000003)	0.000012 (0.000003)	0.000015 (0.000007)	-0.000005 (0.000003)	-0.000019 (0.000003)	0.000034 (0.000003)	-0.000014 (0.000003)
<i>Reset</i> ₃ ²	0.000054 (0.000003)	-0.000001 (0.000003)	0.000011 (0.000006)	-0.000005 (0.000002)	-0.000017 (0.000003)	0.000016 (0.000003)	-0.000003 (0.000003)
<i>Reset</i> ₄ ²	0.000019 (0.000003)	-0.000003 (0.000003)	0.000007 (0.000006)	-0.000006 (0.000002)	-0.000018 (0.000002)	-0.000002 (0.000003)	0.000004 (0.000002)
<i>Reset</i> ₅ ²	0.000025 (0.000003)	0.000013 (0.000003)	0.000004 (0.000007)	-0.000014 (0.000003)	-0.000025 (0.000003)	-0.000015 (0.000003)	-0.000003 (0.000003)
<i>Reset</i> ₆ ²	0.000407 (0.000026)	-0.000224 (0.000026)	0.000229 (0.000053)	0.000047 (0.000022)	0.000061 (0.000022)	-0.000229 (0.000023)	0.000161 (0.000023)
<i>Reset</i> ₇ ²	0.000095 (0.00002)	-0.000452 (0.00002)	-0.000004 (0.000041)	0.000014 (0.000017)	-0.000074 (0.000017)	0.000003 (0.000018)	0.000228 (0.000018)
<i>Reset</i> ₈ ²	0.000340 (0.000235)	0.000600 (0.000238)	0.001190 (0.000484)	0.001004 (0.000198)	-0.004121 (0.000203)	0.000181 (0.000212)	0.000086 (0.000208)
<i>Reset</i> ₉ ²	0.000026 (0.000004)	0.000018 (0.000004)	0.000008 (0.000008)	-0.000019 (0.000003)	-0.000011 (0.000003)	0.000029 (0.000003)	0.000011 (0.000003)
<i>Reset</i> ₁₀ ²	-0.000037 (0.000024)	0.000005 (0.000024)	0.000025 (0.000049)	0.000016 (0.00002)	0.000026 (0.00002)	0.000082 (0.000021)	0.000008 (0.000021)
<i>Reset</i> ₁₁ ²	0.001113 (0.000375)	0.001619 (0.000381)	0.000146 (0.000769)	0.002983 (0.000316)	-0.000315 (0.000324)	-0.000228 (0.000339)	-0.000078 (0.000333)
2016(<i>Indic.</i>)	0.018118 (0.004465)	-0.009527 (0.004531)	0.021905 (0.00924)	0.010483 (0.003763)	-0.003011 (0.003859)	0.039279 (0.004035)	-0.000680 (0.003963)
<i>Constant</i>	0.339842 (0.004247)	0.308457 (0.004678)	0.342375 (0.00966)	0.341581 (0.004395)	0.341538 (0.004294)	0.313380 (0.004389)	0.315620 (0.004242)

Notes: Maximum likelihood estimation results for the discontinuous growth models for all parties, corresponding to the visualizations in Figure 4a and 5 in the main part of this work. All estimates refer to the fixed component (see Equation 1). Standard errors in parentheses.

Table E.3: Within Party Discontinuous Growth Model Results

	AfD	SPD
<i>Time</i>	0.0004 (0.1109)	0.1392 (0.0764)
<i>Time</i> ²	-0.0005 (0.012)	-0.0160 (0.0083)
<i>Event</i> ₁	-0.0418 (0.2675)	0.2072 (0.1767)
<i>Event</i> ₂	0.0682 (0.7055)	2.0552 (1.0742)
<i>Event</i> ₃	0.1759 (1.5794)	8.7424 (4.5511)
<i>Event</i> ₄	0.1690 (3.6542)	25.0055 (12.9016)
<i>Event</i> ₅	0.7395 (15.9836)	62.6707 (32.3396)
<i>Event</i> ₆	1.2022 (27.0927)	94.8064 (48.4635)
<i>Event</i> ₇	1.4875 (31.8545)	104.0130 (53.6669)
<i>Event</i> ₈	2.6329 (38.3479)	117.2648 (60.5441)
<i>Event</i> ₉	2.7083 (39.7186)	120.3558 (61.9692)
<i>Event</i> ₁₀	3.3465 (56.3852)	152.4504 (78.7396)
<i>Event</i> ₁₁	3.6832 (63.1671)	165.4034 (85.335)
<i>Reset</i> ₁	0.0022 (0.4125)	0.1654 (0.1058)
<i>Reset</i> ₂	-0.1088 (0.3044)	0.3764 (0.1953)
<i>Reset</i> ₃	-0.0604 (0.3078)	0.7611 (0.3904)
<i>Reset</i> ₄	0.0237 (0.4223)	1.2772 (0.6549)
<i>Reset</i> ₅	0.0556 (0.8837)	1.9982 (1.0366)
<i>Reset</i> ₆	0.1266 (1.1684)	1.5474 (1.2789)
<i>Reset</i> ₇	0.8230 (1.2509)	2.4962 (1.3395)
<i>Reset</i> ₉	-0.0051 (1.3842)	2.6525 (1.434)
<i>Reset</i> ₁₀	-0.0155 (1.6656)	3.2267 (1.6245)
<i>Reset</i> ₁₁	0.0128 (1.7438)	3.2459 (1.6821)
<i>Reset</i> ₁ ²	0.0422 (0.1913)	0.0140 (0.0144)
<i>Reset</i> ₂ ²	0.0393 (0.0789)	0.0148 (0.0089)
<i>Reset</i> ₃ ²	0.0096 (0.0282)	0.0167 (0.0084)
<i>Reset</i> ₄ ²	0.0004 (0.0121)	0.0156 (0.0083)
<i>Reset</i> ₅ ²	-0.0013 (0.0145)	0.0172 (0.0087)
<i>Reset</i> ₆ ²	-0.0273 (0.0789)	0.2277 (0.0543)
<i>Reset</i> ₇ ²	-0.1432 (0.0434)	0.0397 (0.0299)
<i>Reset</i> ₉ ²	0.0068 (0.0131)	0.0262 (0.0091)
<i>Reset</i> ₁₀ ²	0.0289 (0.0789)	-0.0142 (0.0543)
<i>Reset</i> ₁₁ ²	0.0012 (0.012)	0.0161 (0.0083)
2016(<i>Indic.</i>)	-0.0013 (0.2509)	0.0119 (0.1102)
<i>Constant</i>	0.0094 (0.2174)	-0.1426 (0.1498)

Notes: Maximum likelihood estimation results for the within party discontinuous growth models for AfD and SPD, corresponding to the visualizations in Figure 4b and 4c in the main part of this work. All estimates refer to the fixed component (see Equation 1). Standard errors in parentheses.

Table E.4: Investigating Heterogeneity in Similarity Shifts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	AfD	CDU	CSU	The Greens	FDP	The Left	SPD
Age (60+)	0.00001373 (0.00002506)	-0.00001330 (0.00001987)	0.00003897 (0.00003784)	-0.00000936 (0.00003487)	-0.00000899 (0.00003460)	-0.00001127 (0.00003340)	0.00000920 (0.00003499)
Foreign Population (%)	-0.00001075 (0.00001319)	0.00000198 (0.00000931)	0.00000640 (0.00002039)	-0.00002083 (0.00001679)	-0.00000726 (0.00001718)	-0.00002242 (0.00001666)	-0.00002094 (0.00001906)
Disp. Income	0.00000003 (0.00000003)	0.00000001 (0.00000002)	-0.00000004 (0.00000004)	0.00000006 (0.00000004)	0.00000006 (0.00000004)	0.00000006 (0.00000004)	0.00000008* (0.00000005)
Craftsmen Firms	0.00000739 (0.00003613)	-0.00004803 (0.00003382)	-0.00000962 (0.00003243)	0.00002496 (0.00004980)	-0.00008033 (0.00005621)	0.00001884 (0.00005080)	-0.00009332* (0.00005224)
Unemployment 2017	0.00001516 (0.00003210)	0.00002125 (0.00002540)	-0.00011799 (0.00010342)	0.00005000 (0.00004296)	0.00003005 (0.00004338)	0.00002325 (0.00004464)	0.00003914 (0.00004640)
High Education	0.00000365 (0.00000953)	0.00000077 (0.00000597)	0.00000734 (0.00001350)	0.00000598 (0.00001165)	-0.00000729 (0.00001272)	-0.00000391 (0.00001236)	-0.00000464 (0.00001281)
Manufacturing	-0.00000314 (0.00000505)	-0.00000143 (0.00000378)	-0.00000582 (0.00000496)	-0.00000128 (0.00000651)	-0.00000266 (0.00000670)	-0.00000956 (0.00000636)	-0.00000040 (0.00000700)
East	-0.00025566 (0.00020924)	0.00011190 (0.00014061)	-0.00027317 (0.00026432)	-0.00027317 (0.00026432)	0.00012110 (0.00025564)	-0.00020908 (0.00026373)	0.00004240 (0.00026044)
Constant	9.27056380*** (0.00121043)	-2.27516418*** (0.00061421)	2.20544749*** (0.00079264)	2.27858130*** (0.00141566)	-4.37701452*** (0.00162300)	2.08529032*** (0.00155206)	-8.91505516*** (0.00143820)
Observations	225	189	36	225	225	225	225
R ²	0.023	0.082	0.086	0.028	0.026	0.028	0.038

Notes: Dependent variable is the shift in similarity to the specified parties after the last event. Amount of craftsmen firms per 1,000 inhabitants. Standard errors in parentheses and calculated using bootstrapping. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.