Terrorist Attacks, Cultural Incidents and the Vote for Radical Parties[*]

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

April 12, 2019

Abstract

We study the recent increase in support for a radical right-wing party in Germany (AfD) relating it to the cultural backlash associated with a number of terrorist attacks, which happened in Europe around the same time, and a very salient incident which occurred in the city of Cologne on New Year's Eve 2015/2016. We measure shifts in voters' attitudes after those events by the change in a measure of linguistic distance between the content of the tweets they posted and those posted by the main German parties. Next we use the date of those events to identify exogenous shifts in language similarity at the electoral constituency level, using a random coefficients model. We find that, following such events, the similarity with the AfD tweets increases, while it increases much less or even falls for the other parties. We then regress differences in vote shares between the 2013 and 2017 general elections at the electoral constituency level on exogenous shifts in language similarity, finding that these are significant in explaining changes in vote shares. Our results are a first step towards understanding cultural backlash as a channel of transmission of radical right support, which has been traditionally difficult to study and which we can address in the appropriate way using new data and techniques.

[†]Bocconi University
[‡]New York University

1 Introduction

In recent years in many Western democracies we have observed a shift away from the traditional principles of an open society. The very ideals on which such democracies were founded, including multi-culturalism, have been called into question. Alongside, the salience of immigration and national security in public discourse has increased. These dynamics have moved in parallel with the support for radical right-wing parties, and the corresponding decline of long-established moderate ones.

Common explanations for these phenomena point to economic shocks (the Great Recession) and to globalization, immigration flows and to the increase in inequality, all of which have heightened the malaise of specific groups in society. Globalization and immigration in turn increase the salience of cultural differences and raise the perception of their threats. The distance between cultural norms raises concerns about the quantity of immigrants and increases the demand for stronger borders. Parties which emphasize such threats and promise policies to address them are bound to gain.

We investigate these phenomena studying one specific country and period: Germany between the last two federal elections, that is between 2013 and 2017. In this four year period the German economy was recovering from the Great Recession faster than other Eurozone countries: growth averaged 2 percent between 2014 and 2017 and unemployment fell from 7 percent to below 6 percent although with important differences across regions. Thus in Germany the effects of the economic crisis should have been considerably less salient than in other European countries. Importantly, however, the concerns about immigrants mentioned above were highlighted by two sets of events.

First, a number of terrorist attacks, claimed mainly by the Islamic State of Iraq and the Levant (ISIL), hit various parts of Europe in those years: in Germany there were a total of seven attacks[1], the one on the Berlin Christmas market on December 15, 2016 with a very high death toll: 12 people died and 56 were seriously injured. Other non-terrorist incidents happened in the same period: the most salient was the sexual assault on a number of German women on New Year's eve 2016. In the Cologne city center hundreds of women were victims of sex offences recorded by the police, and similar incidents took place in other German cities: Hamburg, Dortmund, Düsseldorf, and Stuttgart. For all of Germany, police estimated that 1,200 women were sexually assaulted that night and at least 2,000 men were involved.[2] Cologne police chief Wolfgang Albers stated that the perpetrators in his city were reportedly men of "Arab or North African appearance" and said that Germany had never experienced such mass sexual assaults before.[3]

The second remarkable development is the migration of refugees from the Middle East, especially out of Syria and its civil war. In September 2015, Austria and Germany took one of the most pivotal decisions of Europe's refugee crisis, opening their borders to tens of thousands of migrants who had piled up in Hungary. A turning point in the political discourse was a speech by Chancellor Angela Merkel on August 31st 2015 when she said "I put it simply, Germany is a strong country [...] we have managed so many things, we can do this", which was interpreted as an open door to refugees.

At the end of this 4-year period in the September 2017 general elections, Germany experienced a sharp shift to the right. The three parties which had formed a government coalition during the previous legislature suffered big losses: the CDU/CSU and the SPD, the most important center-right and center-left parties of the traditional political system, together lost about 15 percent of

---

[1]Some Jihadi-inspired, not necessarily officially claimed by ISIL, others claimed by ISIL (National Consortium for the Study of Terrorism and Responses to Terrorism 2018).

[2]See Süddeutsche Zeitung (10.07.2016) for a review of the report by the Federal Criminal Police Office of Germany.

[3]See Zeit Online (05.01.2016) for a brief summary.

their vote share and 105 seats out of about 700 members of the Bundestag, the German lower house of parliament. This went to the advantage of the Liberals, whose vote share increased from below 5 to 11 percent, and especially of the far right AfD (*Alternative für Deutschland*) which increased its vote share from 4.7 to 12.6 percent and, importantly, entered the Bundestag for the first time winning a total of 94 seats, 14 more than the liberals. AfD had been created just before the 2013 elections. Initially it was a Eurosceptic party concerned mainly with economic issues; it took a significant right turn after July 2015 following a clash between moderate-right and far-right members. Since then migration and external threats have become its top priorities (Cantoni et al. 2017).

This paper aims to understand what drove the shift in votes observed in the 2017 election. To what extent the salience of issues such as Islamic immigration and external threats affected the election outcome beyond concerns for economic conditions and inequality? We surmise that the reasons for the rise of the AfD should also be found in the cultural shocks associated with the refugee inflow and with the terrorist attacks, as such incidents may have had an influence on German voters. For instance the terrorist attacks occurred in Paris, Nice and London could have influenced the German perception of the threat of Islamic terrorism. The far-right AfD had singled out the inflow of refugees from Islamic countries as a major threat to European identity and security: as violent incidents occurred this party claimed to be proven right. But to what extent did this affect votes?

Because of low frequency and well-known problems of self-reporting, surveys may not be able to isolate the public reaction to these violent incidents. We try to overcome this problem using high frequency data from Twitter and a Natural Language Processing (NLP) algorithm (Doc2Vec) that allows us to compute a metric of distance between the content of the tweets posted by citizens and those posted by the seven main German parties. Following a growing literature in opinion mining (see next section), we surmise that the distance in the use of language captures the alignment between the public mood and the political parties. Tweets are collected whenever produced, so our data is essentially continuous time. Our analysis is at the disaggregated level of German electoral constituencies: this allows us to control for constituency specific characteristics such as the number of immigrants and refugees, unemployment, inequality etc. With an estimated overall size of less than five million Twitter users, Twitter is not the most used social network in Germany. Nevertheless, we argue and discuss that the effect we measure can be seen as a lower bound for a possibly much larger effect in the overall German population.

Language distance could change either because the language used by citizens changes, for example following a terrorist attack, or because what changes is the language of the parties' tweets. In consequence, a decreased distance could result from the public moving closer to the party or vice versa. We show some evidence suggestive that it is indeed the public moving closer the AfD.

Twitter users are not required to reveal where they live: we thus design a rule to attribute them to a specific electoral constituency. To validate our measure of language similarity as a proxy for the mood of voters we analyze about 31 million tweets and show that the measured similarity between these tweets and those of the seven main German political parties is positively correlated with votes and polling data.

Next, we use the date of the events discussed above to identify exogenous shifts in the afore-mentioned language similarity. Our measure of shifts in similarity captures changes in the level or in the trend of similarity and is thus robust to the possibility that in the hours following an event the type of people who tweet might change. We observe that following such events the similarity between the text of the average constituency-level tweets and those issued by the AfD increases, while the change in similarity is lower or negative for the other parties. Changes in similarity

following an event generally differ across constituencies reflecting constituency characteristics. We show that the variation is not explained by standard economic variables, including those who are usually considered explaining the support for right-wing parties. We regress differences in vote shares from 2017 to 2013, at the constituency level, on exogenous shifts in language similarity. We find that shifts in language similarity are significant in explaining changes in vote shares. Hence, we conclude that we found an exogenous factor beyond the economic variables research usually considers, explaining electoral outcomes, from which the right-wing party AfD benefited the most.

We then perform placebo checks to ensure that comparable language changes are not observed during other salient public events. Finally, we discuss the objection that voters are becoming more similar to AfD and not the other way round, which would also cause a lower distance. We also address the possibility that non-cultural issues on the political agenda of the AfD, such as its negative stance towards the euro, are driving the increased similarity, concluding that this is unlikely to be the case.

2 Related Literature

Economic Determinants of the Growth of Populism

Our work intersects several strands of literature. We speak mostly to the area of political economy trying to explain the rise of populism in Western democracies. Although this literature has grown fast in recent years,[4] the study of modern far-right political formations dates back to at least the 1990s (Betz (1994) and Kitschelt (1995)).

The rise of populism is associated to rising unemployment and increasing economic insecurity (Knigge (1998), Funke et al. (2015), Inglehart et al. (2016), and Rodrik (2018)). Colantone et al. (2018c), Colantone et al. (2018a), and Autor et al. (2016) establish a clear connection between displacements due to globalization and import competition and the rise of nationalist parties and of politicians advancing protectionist views.

Some studies exploiting micro data suggest different explanations for right-wing populist support, among which protection of national identity against outsiders (for instance, Oesch (2008)). More recently, using a vast collection of administrative data, Dal Bó et al. (2018) argue that populist voters and candidates in Sweden share some characteristics of social groups affected by increased insecurity in the labor market.

This literature also highlights the interaction between economic shocks and socio-cultural attitudes, among which the decline of trust in institutions. For example, Algan et al. (2017) find that lower trust in political institutions is associated with a rise in unemployment and a fall in GDP. Similar findings are described in Dustmann, Eichengreen, et al. (2017).

Immigration and the Growth of Populism

Another reason for populist support often discussed is reaction to immigration. While the traditional literature has long identified a correlation between populist vote and anti-immigration attitudes (for instance Golder (2003)), recent work has focused on identifying the causal effects of immigration flows and disentangling heterogeneous channels. Using data from Italy, Barone et al. (2016) show that immigration moves votes to the right of the political spectrum and that the main drivers are fears of increased competition in the labor market and for welfare. Dustmann, Vasiljeva, et al. (forthcoming) study a Danish policy of random refugee allocation and find that

---

[4]See Arzheimer (2019) for a comprehensive bibliography.

more refugees increase votes for right-wing parties, the effect being larger in smaller and more rural municipalities and in those with higher crime rates. Competition for compositional amenities, that is the quality of local public goods and services, is another identified channel, starting from the work by Card et al. (2012). Halla et al. (2017) find that immigrant inflows in Austria induce fear of deterioration of local services and move voters towards the far-right. Hangartner et al. (2017) come to even more extreme conclusions: studying the case of Greek islands, they show that even the exposure per se to refugees increases native hostility towards immigrants and minorities. Somehow mixed findings are presented in Steinmayr (2016) for Austria and Vertier et al. (2018) for France, who show how exposure to refugee relocation can in some settings decrease support for the far-right. Notably, in some cases the perceived threat from immigration can exceed the actual threat by large amounts as found, for example, by Alesina et al. (2018) using survey data and experiments in six countries.

As mentioned, Germany seems a particularly suitable context in which to study different drivers of populism other than economic distress, since the country was much less affected by the Great Recession compared to many other Western countries. It is interesting to note that for the party we are focusing our research on, AfD, some scholars find clear differences relative to other Western European right-wing parties (Hansen et al. (2019), Goerres et al. (2018), as we shall discuss in more detail later.

Twitter, Attitudes and Politics

Our work speaks also to the literature on social media, Twitter in particular, and political behavior. Due to its easily accessible Application Programming Interface (API) and its grown popularity, data from Twitter are increasingly used to investigate a variety of topics in social sciences. Differently from surveys administrated through questionnaires, data from social media allows to "nowcast" attitudes, as noted by Curini et al. (2015), including tracing their fluctuations after the occurrence of specific events. Another advantage of using Twitter, compared to other social networks, is its nature of openly exchanging messages ("Tweets") as either a statement on its own or as a reply to another message. Also, the ease with which common topics of discussion ("Hashtags") and the large body of information incorporated in each Tweet, including data about the author and its networks, can be retrieved from the API makes Twitter attractive to researchers. The use of Twitter data to the study of politics include, among others, investigations of the political preferences of either a population of individuals (Barberá et al. (2015), Huberty (2015), and Vaccari, Valeriani, Barberá, Bonneau, et al. (2013)) or more specific actors of public interest (An et al. (2012), Park et al. (2015), Cantoni et al. (2017), and Sterling et al. (2018)). Other topics for which Twitter provides a valuable resource are the diffusion of information (Halberstam et al. (2016), Vaccari, Valeriani, Barberá, Jost, et al. (2016), Gorodnichenko et al. (2018), and Jann et al. (2018)) and the use of networks in the context of political campaigns (Safiullah et al. (2017), Huberty (2015), and Jin et al. (2017)).

Text Analysis

In this paper we use a Natural Language Processing (NLP) algorithm - Doc2Vec - to learn a vectorial representation of the language used in Tweets. The last decade has seen a rapid evolution of applications of text analysis. For an overview of applications connected to opinion mining, see Liu (2012) and Ravi et al. (2015). Researchers tend to use these techniques to either analyze the language of one group of interest with respect to some pre-defined characteristics (Curini et al.

(2015), Sterling et al. (2018)) or compare languages of two or more groups of interest (Gentzkow et al. (2010), Obschonka et al. (2018)).

Language, Terrorism and German Politics

Our paper is also closely related to recent work on the roots of support for the German far-right, the effects of terrorism on attitudes and partisan differences the use of language in social media. Among the first category, Cantoni et al. (2017) use text analysis and Twitter data to measure the persistence of right-wing ideology in Germany. Interestingly, they show the existence of a historical correlation between support for the Nazi party in the 1920s and 1930s and support for the AfD from 2015 onward. Among the second, though not in the German context, Sterling et al. (2018) compare differences in language use between liberal and conservative U.S. residents and members of the U.S. congress. They focus on values implicit in the tweets of both groups, such as group loyalty and authority, and find that indeed language use reflects characteristics which are commonly associated with these groups, for example that liberals are more concerned about fairness issues. Finally, some recent work has studied public opinion in relation to the recent wave of terrorism. Giani (2017) uses the coincidence of the collection of a social survey at the time of the Charlie Hebdo shooting in Paris in January 2015 and finds an increasing support for institutions engaged in counter-terrorism, following the attack. More surprisingly, his findings show a decreasing support for reducing Muslim immigration. Similarly, Legewie (2013) exploits terror attacks in Indonesia and Madrid as natural experiments to analyze the effect of terrorist attacks on attitudes toward immigrants in European countries using Eurobarometer data. He finds evidence of cross-national variation, which is correlated with local characteristics. For example, increasing unemployment rates are positively correlated with the size of the effect of those attacks.

3 Background: AfD, Terrorism and the Refugee Crisis

The Rise of AfD in German Politics

AfD (*Alternative für Deutschland*), was founded at the beginning of 2013 by an academic economist, a conservative politician and a journalist. The new party aimed at promoting economic and fiscal conservatism and giving representation to groups of civil society which were worried by the potential risks that the European mechanisms of assistance to distressed countries in the South of the continent could pose to German public finances.

AfD first participated at the 2013 Federal Elections, barely missing the 5% threshold of nationwide votes for winning seats in the German Parliament.[5] Still, its first-time performance was remarkable.[6] The party gained a large part of its votes in the Eastern regions of the country.

At the European elections of the following year, AfD ran again with a Eurosceptic program and gained 7.1% of nationwide votes, securing its first seats in European Parliament. By that time, a strong tension had emerged within the party between the more centrist, fiscally conservative faction led by economist Bernd Lucke and a grown far-right leaning wing, markedly nationalistic and anti-immigration, led by an entrepreneur Frauke Petry. As AfD rose in its electoral successes, winning seats in regional elections and consolidating a constituency in the Eastern regions of the country, the national conservative faction gained importance. July 2015 is the moment which political scientists and journalists see as the crucial point in the life of AfD. A national convention elected Petry as

---

[5]In the German electoral system, voters cast two votes: one for a constituency representative elected through a first-pass-the-post system and one for party lists.

[6]See for instance Spiegel Online (23.09.2013).

party leader while Lucke left the party with thousands of members of the moderate wing to later create another political formation. From that moment, the party agenda shifted markedly from economic issues to criticism towards the immigration policy brought about by German Chancellor Angela Merkel (and, soon after the 2015 refugee crisis overlapped with a season of Islamist terrorism in Europe) to the connection between multiculturalism and threats to security. The content of the party's messages in the social media reflect this tendency, with nationalism and Islam rapidly replacing the Euro among the main topics in the AfD accounts (Cantoni et al. 2017).

Between 2015 and 2017 AfD ran several state elections, which confirmed its growth and elected representatives in all state Parliaments, at the expenses of the more moderate center and of left-wing parties, most notably the CDU and the SPD. At the 2017 Federal election, AfD was the third party in terms of votes in the nationwide list and of seats gained in Bundestag, the German Parliament.

In recent years, after the party turned into a far-right populist formation, political scientists have studied its electoral base. The most recent cross-country evidence on populist voting has underlined the importance of "globalization losers" in the support of populist parties in Europe (Colantone et al. 2018b). Case studies of AfD have shown that attitudes about immigration policy and German institutions (Hansen et al. 2019), as well as conservative economic policy preferences (Goerres et al. 2018) rather than economic insecurity are at the basis of support for the party. Cantoni et al. (2017), as already mentioned, bring evidence in favor of cultural channels behind populist voting in Germany: they show that after the nationalist turn of 2015 the geographic distribution of votes for AfD has become correlated with that for the Nazi Party.

The Refugee Crisis and Terrorism

The social and political context in which AfD rose to power was marked by a dramatic increase in immigrant flows to European countries, including Germany (see Figure 1). Major military conflicts in Syria and Sub-Saharan Africa, together with economic factors, generated large movements of people towards North and Western Europe, through the Mediterranean Sea or the Balkans. According to Eurostat, almost 4 million people applied for asylum in one of the European Union countries from 2014 to 2017. The political reactions in the Continent were not uniform. German Chancellor Angela Merkel announced in September 2015 that she would open German borders for Syrian asylum-seekers, thus inaugurating an open and welcoming policy towards refugees. The refugee policy became soon one of the crucial issues around which AfD concentrated its opposition to the Merkel government.

Parallel to the refugee crisis, starting in 2014, Western Europe was hit by an unprecedented wave of religious terrorism. A number of terrorist attacks were perpetrated across Europe, including France, Germany, Belgium and the United Kingdom.[7] Some of the perpetrators were affiliated to ISIL or to other radical Islamist organizations, most notably Al Qaeda, but in several cases the attackers were common citizens, often second-generation immigrants, Muslim or recently converted, which had rapidly radicalized. While most of the attacks were carried out with weapons, guns and bombs, "lone wolf" terrorists often resorted to common objects which were immediately available to inexperienced killers and which were readily converted in instruments of death: knives, trucks and autovehicles. Some of the deadliest attacks of that period were actually executed with such improvized weapons. The unexpected violence of the attacks and the large number of victims, together with their relatively high frequency (from 10 to 15 between 2014 and 2016, according

---

[7]The first violent attack of this upsurge of Islamist terror was the shooting which happened in January 2015 at the headquarter of magazine *Charlie Hebdo* in Paris and, few days after, in a kosher supermarket of the same city.
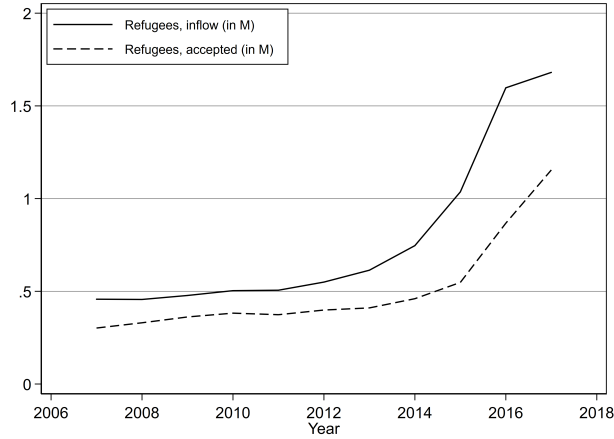
Figure 1: Refugee Inflow and Legal Acceptance

to different classifications) and the new profile of "next-door terrorist" which emerged, generated widespread fear and a sense of threat in the public across Europe. The political debate about security and immigration policy was inflamed and far-right and nationalist parties gave voice to fears about the threat coming from multiculturalism and failed assimilation of immigrants. In Germany, criticisms about Merkel's government and its immigration policy grew in light of these events.[8]

This wave of terrorist attacks - as we shall explain in the following section, there was also another salient cultural incident in Cologne on New Year's eve 2015/2016, not related to terrorism - opened the door to a cultural backlash against immigrants. One of the main drivers was the perception of an external cultural threat which became abruptly more salient. In the next sections we use some of these terror events to derive measures of shocks to attitudes about culture.

4 Data

The data used in this paper includes data from German Twitter users, statistics on municipalities and refugees provided by the Federal Statistical Office of Germany, and statistics on macro variables and electoral results at the constituency level provided by the Federal Returning Officer. We also use polling data at state level provided by infratest dimap (2018). The use of the different sources is described in the following section.

Twitter Users

The goal was to build a sample of German Twitter users which is representative of voters in German electoral constituencies. We started from a complete list of municipalities belonging to each constituency provided by the Federal Returning Officer (Der Bundeswahlleiter 2017). The first challenge was to identify where Twitter users live, i.e. the municipality where they are registered to vote. Twitter users can voluntarily choose to publish any location they wish on their profile and there is no mechanism in place to double check the provided information. Hence, using information provided by users leads to four possible outcomes. Missing addresses, correct addresses, incorrect addresses, and non-existing addresses (such as Disneyland). Excluding the latter is straightforward, but there does not exist a simple method to verify a given existing address or to replace a missing one. Most importantly, it is not possible to verify whether the location that an user provides is her

---

[8]See for instance Politico (20.12.2016).

(a) Case 1   (b) Case 2

Figure 2: Sampling Mechanism

real place of residency or not. For this reason, we constructed an algorithm that allocates users to a constituency even if they do not provide their location.

The algorithm works as follows. The 299 German electoral constituencies are drawn with the goal to equalize population across them. In consequence, electoral borders in general do not follow any common structure, but are drawn over cities and districts.[9] By the end of 2017 there existed 401 districts and district-free cities, which correspond to the NUTS-3 classification of the European Union (see European Parliament and the Council (2003)). For a given constituency our approach first identifies the largest towns within each district that is inside the given constituency. Here we face two possible situations (as shown in Figure 2).

For a constituency (such as C_1) which 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 and T_2). Because one district (D_2) can overlap with several constituencies, the chosen towns are not necessarily the largest towns in their districts (T_3 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 - 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. In the end, our sample consists of 261 constituencies, either original or artificially merged, in which the rule described above produces a sample of 493 towns.

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, theaters and more. Other accounts, such as sport clubs, TV stations or newspapers were not considered, because their attraction is not bound to local residents. For example, we assume that following a famous football club or a well-established newspaper is not a reliable source to infer where a users 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.

Having identified local landmarks, we used the Twitter API to fetch all the followers of those landmarks. We then restricted the sample, dropping users who followed less than three landmarks in the same constituency or landmarks in more than one constituency. This strategy produced a fi-

---

[9]They also change from one election to another. The number 299 corresponds to the general election of 2017 (Der Bundeswahlleiter 2017).

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
ΔAfD Vote (Mean)	0.0809	0.0981	0.0832
N	225	36	261

Note: Based on own calculations, including merge of large cities. In total we have 261 constituencies in our sample, for which we obtain continuous Twitter data for 225. ΔAfD Vote refers to the difference in vote share from 2013 to 2017.

nal sample of 178,271 Twitter users. The exact amount of active German Twitter users is unknown; different sources estimate it between 2 and 5 million users.[10] For our sample we downloaded all available tweets. We do not consider retweets because we want to avoid having text we did not allocate to a constituency. Also, assuming that automized Twitter users are unlikely to follow local landmarks, not considering retweets limits the threat of including text not written by actual users ("bots"). Twitter limits the access to roughly the latest 3,200 tweets, but since this limit was only reached by 128 users we consider its influence on our research negligible and conclude that we have essentially all the tweets sent by our sample of users.

Possible Sources of Bias in the Data

We follow the strategy of including in our sample Twitter users from all German constituencies. This could give rise to two different form of biases. First, the constituencies for which we cannot observe any users, due to a lack of activity in the sampled cities, might be different with respect to support for the AfD, compared to the ones we observe. Table 1 shows the mean vote share for the AfD in the constituencies we observe and in those we do not observe. Note that these shares are based on our own calculations after merging large cities and only reflect the vote of the party-specific vote. We observe that in the constituencies we sample, the AfD vote share is slightly lower than in the ones we do not observe. This presents evidence that our sample of constituencies does not consist of more-than-average right-wing supporters.

The second possible source of bias lies within constituencies. We have no reason to believe that a representative Twitter user from our sample equals a representative German voter in a given electoral constituency. There is clearly a mechanism of self-selection into our sample. However, we find two pieces of evidence suggesting 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).[11] 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 are older were registered on the social network (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.

Of course, this evidence is not a proof rejecting 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

---

[10]In 2016, Twitter announced that 12 million people are in touch with Twitter in Germany, including both registered active users and non-registered users. Considering only registered users will lead to a (much) lower number, which can only be estimated.

[11]The 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. Detailed information can be obtained from The Federal Returning Officer.

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

Retrieved February 11, 2019. Amount of tweets includes retweets to better reflect activity.

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 *Bündnis 90/ Die Grünen* issued roughly the same amount of tweets and retweets (although over a longer period of time), there are more than three times as many users following this party, compared to the amount of followers for the AfD. In fact, the AfD is the party with the fewest amount of followers on Twitter, although it exceeded three of those parties in vote share. This represents strong evidence that Twitter users are not overly supportive of the AfD.

Thus the demographic profile of Twitter users in Germany is not close to that of a representative supporter of a right-wing populist party.[12] Hence, a representative user from our sample is likely to be more moderate or liberal in political beliefs than a potential AfD voter. Therefore, we argue that although we will not be able to identify an externally valid effect for the German electorate, we will probably underestimate it. Simply formulated, we argue that if there is an effect in the population of Twitter users, then one has reason to believe that there is even a stronger effect in the German population.

Parties

The parties considered in our research are those which made it into the *Bundestag* in 2017: *Alternative für Deutschland* (AfD), *BÜNDNIS 90/DIE GRÜNEN* (Die Grünen), *Christlich Demokratische Union Deutschlands* (CDU), *Christlich-Soziale Union in Bayern* (CSU), *Die Linke*, *Freie Demokratische Partei* (FDP), and *Sozialdemokratische Partei Deutschlands* (SPD). For each party we considered the main national level Twitter account, excluding the personal accounts of party leaders and representatives.[13] Table 2 shows descriptive statistics of those accounts. To overcome Twitter API limits, which allow to download only the latest 3,200 tweets, while these party accounts differently from our sample of users all have many more tweets, we web-scraped the parties' Twitter pages directly.

Electoral and Structural Data

For the years of federal elections, the Federal Returning Officer in Germany (*Bundeswahlleiter*) publishes the election results and structural variables for each electoral constituency, since, as we have seen, electoral constituencies do not follow the borders of administrative districts (NUTS-3). Note that for all electoral analysis we only consider the second vote (that is the vote for the party list), since it is the most important for the distribution of seats in the German Bundestag and

---

[12] This claim is also supported by Roth et al. (05.10.2017).

[13] The reason is that we want to assess comparable accounts for all parties. Some party leaders, such as Angela Merkel, are not present on Twitter. Including party representatives requires subjective assessment of this set of people, which we want to avoid.

represents purely party preferences, free of possible preferences for a local candidate. Data on the amount of refugees within each electoral constituency were downloaded at the administrative district level from the Federal Statistical Office for Germany and aggregated at the constituency level. For the constituencies we merged due to their size, a weighted average with respect to population was used to aggregate variables of interest.

Apart from electoral results we also use polling data at state level provided by infratest dimap (2018). These polls are the results of the so called "Sunday's Question" (*Sonntagsfrage*): more than 1,000 eligible voters are asked which party they would vote for if there were a General Election next Sunday. As indicated by the institute itself the answers to this question should be treated carefully, since they reflect the current mood of the electorate rather than a precise forecast of voting behavior. For us, however, this offers a chance to test our measured index of similarity. In the end, we are not claiming that what we measure can be used to perfectly forecast an election months or years later. What we claim is that it reflects the alignment of political views to a given party.

Events

For our analysis we use eleven events, from the end of 2015 until close to the parliamentary election in 2017. These events were chosen because they represented large shocks to public opinion. Ten out of eleven events are major terrorist attacks. Among the several events related to terrorism and violence 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 which were independent of a national context (theft 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. city or district), but in the whole country. Thus, we exclude some non-deadly attacks and relatively less important events. Third, we select events which plausibly affect 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 unfortunately realistic. The ten terrorist attacks that we consider are:

- *November 13, 2015 in Paris, France.* Three groups of terrorists perpetrated simultaneous attacks with bombs and guns on several targets in Paris, including a stadium, restaurants and a concert hall. 130 civilians were killed and other around 400 injured.

- *March 22, 2016 in Brussels, Belgium.* Three coordinated bombings at several locations in Brussels, 32 people were killed and more than 300 injured.

- *July 14, 2016 in Nice, France.* A young man drove a heavy truck at high speed over the crowd in a boulevard in Niece. 86 civilians were killed and around 400 injured.

- *December 19, 2016 in Berlin, Germany.* A young man drove a truck over the crowd in a Christmas market in the German capital city. 12 civilians were killed and around 50 injured.

- *March 22, 2017 in London, United Kingdom.* A man drove a car into pedestrians in London, injured 50 and killed 5 people.

- *April 20, 2017 in Paris, France.* An attacker shot three policemen and another person, killed three of them and injured the fourth one.

- *May 22, 2017 in Manchester, United Kingdom.* A suicide bombing after a concert at Manchester Arena. 22 people were killed and more than 500 injured.

- *June 3, 2017 in London, United Kingdom.* A car drove into pedestrians in London. 8 people were killed and 48 injured.

- *August 16, 2017 in Barcelona, Spain.* Bombs detonated and a car drove into pedestrians in Barcelona. 16 people were killed, over 100 injured.

- *September 15, 2017 in London, United Kingdom.* A bomb detonated at a train station, 30 people were injured.

In addition to those, we include a non-terrorist event which shocked public opinion in Germany and across Europe and which generated wide political and social reactions consistent with the cultural threat framework.[14]

- *December 31, 2015 and January 1, 2016 in Cologne, Germany.* During the New Year's Eve celebrations in the Cologne city center (and, to a smaller extent, in other German cities), several hundreds of women were subject to mass harassments and sexual assaults, often resulting in thefts. Most reports to the police identified North African and Syian young men as the perpetrators.

5 Similarity between Texts

We compute similarity between the Tweets of the parties and the tweets of each constituency by transforming the two groups of tweets into vectors through Doc2Vec. The similarity is then measured by the cosine similarity between the two vectors. Before proceeding with Doc2Vec, we pre-process tweets.

Text Preprocessing

All the words in our corpus of tweets represent the vocabulary upon which our textual analysis is conducted. Since computational time is more than directly proportional to vocabulary size, text-preprocessing is the practice of reducing the initial vocabulary in such a way that relevant information is not lost but computational time is greatly reduced. We followed standard procedures in text pre-processing with different libraries in Python. We lower-case all words and tokenize the text, i.e., we break streams of text into single words, called tokens. We use word_tokenize from the Python module NLTK. Next, we eliminate punctuation and stopwords, namely words that recur very frequently in our corpus and, thus, have little meaning. The dictionary of stopwords is the one in NLTK. We also remove all tokens that consist only of non-alphanumeric characters; for tweets we 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". This stemming relies on pre-existing dictionaries. We use the German Stemmer in the Python module gensim. Finally, we perform collocations, namely, identifying combinations of two words that have higher probabilities of occurring together than separately. For instance, the tokens "angela" and "merkel" have higher chances of co-occuring as the bigram "angela merkel" than separately. In this case, collocations transform the two separate tokens into just one: "angela_merkel". We

---

[14]See for instance Reuters (10.01.2016).

used BigramCollocationFinder in NLTK. The next step is to learn how to represent these words as vectors.

Doc2Vec

After pre-processing our tweets, we create two types of "documents": a party- and a constituency-document. A party document is the text of all the tweets that the party posted a certain day. A constituency document is the text of all the tweets that all the users in our sample located in a certain constituency posted a certain day. Since we have 750 days in our observation period (from September 4th 2015 – the day of Angela Merkel's speech – to September 24th, 2018 – the day of the election), we end up with 750 documents for each party and 720 documents for each constituency in our sample.

We rely on 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).

As an extension of Word2Vec, Le et al. (2014) introduced Doc2Vec to learn embeddings of documents. Doc2Vec is an extension of Word2Vec that learns to capture not just individual words, but entire sentences and paragraphs. 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: the algorithm is instructed to train the Doc2Vec model several times with unlabeled data while exchanging the input order of the documents. 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.

- Subsampling: The threshold for configuring which higher-frequency words are randomly downsampled, useful range is (0, 1e-5). We set it to 1e-3.

- Negative: The number of "noise words" that should be drawn. We set it to 5.

14

Computing Similarity

The Doc2Vec gives us vectors representing each party and each constituency documents. We then measure similarity between party p and constituency c in day d as the cosine similarity between the two corresponding vectors (i.e., the normalized inner product of the two vectors):

$$\cos \theta_{cpd} = \frac{\overrightarrow{c_d}\,\overrightarrow{p_d}}{\left\|\overrightarrow{c_d}\right\| \left\|\overrightarrow{p_d}\right\|}$$

6 Validation

The natural question arising after applying our algorithm is whether the computed similarity is indeed a valid measurement for how close public opinion is to the various national parties. For the purpose of validation we need data of public opinion for comparison. As a first check we use the electoral results in which we are most interested in, i.e. results for the 2017 federal election at constituency level to test our measure of similarity. Furthermore we use results of poll surveys provided by infratest dimap (2018) at state level as a secondary test. In practice, we investigate whether the index of similarity is positively correlated with both measures.

For the election outcomes, we merge the texts we obtain from Tweets in the 30 days before the election within electoral constituencies and apply the doc2vec algorithm. The same is done for 15 days before the election as a robustness check. The reason for this merge is to produce a considerable amount of text for both parties and constituencies: not all parties issued Tweets close to the election. This may seem counter intuitive, however note that we are only using Tweets, i.e. messages written by the main party account itself, and not any other kind Twitter activity, such as Retweets or replies. For the analysis of poll data, since poll surveys are conducted at state level, we merge the texts from Tweets of all constituencies within a given state on the day of the poll.

We then use as dependent variable the change in vote share from 2013 to 2017 for all parties combined[15] and the poll results obtained and regress them on the measured similarity. We cluster standard errors on the lowest aggregate for the units of observation, i.e. electoral constituency level and state level. 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 3. We observe a positive correlation in all three columns. We conclude that our computed similarity captures the public mood across states and electoral borders to a degree that enables us to move on.

7 Identification Strategy

Effect of Events

We are interested in identifying the effect of the set of events presented in section 4 on the support for political parties at the electoral constituency level. Our analysis relies on the plausible assumption that this subset of events represents exogenous cultural shocks whose occurrence is independent of any local conditions. The size of the 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 random.

Recall that we measure similarity to all parties across German electoral constituencies and over time, resulting in a panel data structure. In this panel the frequency of the data is daily.

---

[15]That is, each observation is party-constituency or party-constituency-date.

Table 3: Validation: Measured Similarity

	ΔVote Share		Poll
	(1)	(2)	(3)
Similarity			
15 days before Election	0.0982***		
	(0.0164)		
30 days before Election		0.126***	
		(0.0186)	
2015 to 2017			0.0679***
			(0.0218)
N	1122	1200	454
R^2	0.020	0.031	0.684

Notes: ΔVote Share refers to the difference in electoral results between 2017 and 2013. The independent similarity is the measured similarity in the specified period. We merge text 15 and 30 days before the election. For polls, it corresponds to the day the poll was conducted. In estimations of electoral outcomes (1) and (2) standard errors are clustered on constituency level. Poll refers to poll surveys on state level conducted between 2015 and 2017. In estimation of poll results (3) standard errors are clustered on state level and party fixed effects are included. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

One way to study the effect of events on similarity is to compute the difference between the similarity prior to the event and the one after it happened. However, inference based on this value faces 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. Moreover, we could simply measure an immediate outrage, while what we are interested in is the long-run effect of those events on the language similarity between people and parties. That is, we want to investigate whether there exists a positive or negative shift towards parties that can be attributed to those events.

To measure this long-term shift in similarity we use the so called discontinuous growth model.[16] This model allows, at specified points in time, for a change in the slope and intercept. In other words, after each event both the time trend and the level of similarity to parties are allowed to shift. The estimated coefficients will be relative to an estimated trend in absence of any discontinuity. In this way we move beyond the immediate reaction to an event and capture its long-term effect on the evolution of similarity over time.

Exploiting the panel structure, we are able to estimate shifts in similarity at the electoral constituency level using a random coefficients model. To check whether there is variation across constituencies, we run an unconditional means model and we estimated the proportion of total variance that occurs between constituencies. Since we find it to be about ten percent, we argue that there is enough variance to justify the use of random coefficients. That means, we allow for changes in intercept and time trend of similarity in each constituency at the day of events. For each party p we estimate separately

$$
\begin{aligned}
simil_{ti}^p = \pi_{0i}^p + \pi_{1i}^p Time_t + \pi_{2i}^p Time_t^2 + \mathbf{E_t}\pi_{3i}^\mathbf{p} \\
+ \mathbf{Reset_t}\pi_{4i}^\mathbf{p} + \mathbf{Reset_t^2}\pi_{5}^\mathbf{p} + \pi_6 Year + \epsilon_{ti}^p
\end{aligned}
\tag{1}
$$

where the index i in the coefficients represents the respective random coefficient. For example, $\pi_{1i}^p = \pi_1^p + r_{1i}$ where r_{1i} represents the random coefficient allowing for between-constituencies

---

[16]The discontinuous growth model is also referred to as piecewise hierarchical model or multiphase mixed-effects model. See Bliese et al. (2016) for an overview of the usage of this model.

differences in time trend.[17] $simil_t^p i$ is the measured daily similarity for party p in constituency i over all periods; $Time_t$ and $Time_t^2$ are the time and quadratic time trend and estimate how similarity would evolve in the absence of events. $\mathbf{E_t}$ is the matrix containing the event indicator vectors, coded 0 prior to the event and 1 after the event: the associated vector of parameters $\pi_{\mathbf{3i}}^{\mathbf{P}}$ estimates the extent to which the predicted value of this model on the day of an event differs from the predicted value in absence of *any* event, and which is based on the trend prior to the first event. $\mathbf{Reset_t}$ and $\mathbf{Reset_t^2}$ are matrices containing event-specific variables coded 0 until the day an event first occurs and then increasing day after day.[18] The associated set of parameters $\pi_{\mathbf{4i}}^{\mathbf{P}}$ indicates the degree to which the event alters the slope π_{1i}^p for the linear effect of time within constituencies, while the set of parameters $\pi_{\mathbf{5}}^{\mathbf{P}}$ indicates the extent to which the event alters the quadratic effect of time estimated by π_2^p. $Year$ is a standard year fixed effect equal to 1 in 2016 and 0 elsewhere.[19]

Note that the matrix $\mathbf{E_t}$ contains 11 event indicators. Hence, for $k = 1, \ldots, 11$ we estimate coefficients $\{\pi_{3ip}^k\}$ which represent the accumulated shifts in similarity to all parties across electoral constituencies after an event occurred. Notice that these shifts cumulated across events. That means the shift after the last event is based on the predicted value of the model which is incorporating all prior events. In case the events caused a downward shift in similarity to parties, the coefficients show a negative sign (and vice versa).

Next, we pool all parties together and estimate the average effect of changes in similarity induced by the complete set of events on electoral outcomes. This overall change is represented by the coefficients of the eleventh event. Define the vector $\Pi \equiv \{\pi_{3ip}^{11}\}$, we estimate

$$\Delta vote = \alpha + \beta\,\Pi + \upsilon \tag{2}$$

where the vector $\Delta vote$ contains the difference in vote share for each party across constituencies between the general elections of 2017 and 2013. Differently from papers which explain party votes with economic variables such as unemployment, we relate votes to language similarity. Note that, different from variables such as unemployment, which have a single value (for each party) for a constituency, our right-hand side variable can vary across parties within a constituency. Hence, whereas using macroeconomic variables as independent variables is not valid when pooling all parties together (because the independent variable does not vary within a constituency, while the dependent variable does) Π is.

Heterogeneity Across Constituencies in Similarity Shifts

As described above, although events occur independently of local characteristics, their effect on similarity is likely to be determined by local conditions. What constituency characteristics determine variation in similarity shifts across constituencies after an event? In particular, we are interested in understanding what constituency level characteristics can explain the variation in shifts towards or away from the AfD. We address this question starting from the estimated coefficients for the last event, i.e. the difference between the predicted value of our model and the predicted value in the absence of the discontinuities caused events, and regress it on a set of local characteristics provided at constituency level. The choice of this set of variables is difficult, since it is hard to guess what kind of variables amplify or reduce the measured reaction to terrorist attacks. We have limited our analysis to investigating whether the standard variables considered for explaining the growth of populist parties, such as unemployment, the share of employees working

---

[17]We were not able with our data to allow for random coefficients on squared terms.
[18]See Table 6 in the Appendix for an overview on the coding of the variables.
[19]We include only one year fixed effect due to high multicollinearity in our estimation.

in manufacturing or foreign population, can explain the variation across constituencies in shifts of similarity. We thus estimate for each party p

$$\Pi_i^p = \alpha^p + \mathbf{Covariates_i}\,\beta^{\mathbf{p}} + \eta_i^p \tag{3}$$

where $\mathbf{Covariates_i}$ is a set of macroeconomic variables at electoral constituency level and Π_i^p represents the vector of cumulated positive and negative shifts in similarity to the parties after the last event. Our special focus lies on the AfD, but for the sake of completeness we shall also report the results for other parties.

8 Results

Cumulated Shifts in Similarity

Results are presented in Figure 3. For $k = 1, \ldots, 11$ and different parties p we show the distribution across constituencies of the coefficients π_{3i}^{kp} representing the difference between the predicted level in absence of events and the predicted value produced by our model which incorporates discontinuities. The parties shown are the AfD and for comparison the traditional center-left SPD. The exact model we estimated using maximum likelihood is reported in Equation 4 in the appendix. Note that, although this graph presents boxplots (in which the horizontal lines represent the median change in intercept across electoral constituencies) variation across constituencies appears to be negligible.

Changes in language similarity at events is positive for the AfD and negative for the SPD. This effect cumulates over time: the final event happened two weeks before the election.

Results for the other parties are presented in Figure 4 and are in line with what we would have expected. Compared to the AfD, none of the other parties show an equivalent increase in predicted similarity.

Interestingly, the shifts in party similarity do not follow precisely the anti-immigration stances of parties. For example, the CSU, the Bavarian ally of 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 in line with recent party history: the union of CDU and CSU was under enormous pressure during the peak of the refugee crisis around 2015. After Angela Merkel openly declared an open-border policy for asylum seekers she was challenged by high-ranked CSU officials who questioned her leadership by promoting closed borders and deportation.[20] Thus, observing a positive shift in language similarity for this party and the AfD is not surprising.

Note however, note that being a left party, such as *Die Linke* or *Die Grünen*, does not necessarily imply negative changes in similarity, which is somewhat counter intuitive. This observation motivates further investigation into the exact words used in tweets. Further work should investigate on what the public and the parties are talking about in the hours after our events. We hope that this can give insights in the nature of changes in similarity.

Cumulated Shifts in Similarity and Votes

The last section has shown how terror events affect the language distance from right-wing parties relative to center and left parties. A natural question arising is if this change in similarity is able to explain electoral outcomes in the 2017 federal election: more specifically the change in each party's 2017 vote share relative to 2013. Our research design enables us to proceed without controlling

---

[20]See Foreign Policy (22.10.2015).

(a) AfD

(b) SPD

Figure 3: Cumulated Shifts in Similarity, AfD and SPD

for covariates due to the use of the random coefficients model presented in Equation 1: that is, we know that certain time-varying and time-invariant differences between constituencies affect the (movement) of similarity to parties — and for this reason we allow electoral constituencies to have their own trend over time and intercept. Events, however, are exogenous to local conditions. Furthermore we are not trying to assess which local characteristics explain electoral outcomes in general: what we want to investigate is whether cultural incidents and terror events have independent explanatory power for electoral outcomes.

Results are presented in Table 4. The dependent variable in this regression 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 Π_{ip} across parties and constituencies. We observe a positive and significant relationship between shifts in similarity and electoral outcomes. After the large differences presented in Figure 3 and Figure 4, where the AfD appears to be the party with the strongest upward shift, this should not be a surprise considering that the AfD was the party

19

(a) CDU

(b) CSU

(c) Die Grünen

(d) Die Linke

(e) FDP

Figure 4: Cumulated Shifts in Similarity, other Parties

Table 4: Electoral Effect: Votes on Cumulated Shifts in Similarity

	Δ**Vote Share**
Cumulated Shifts in Similarity	0.01***
	(0.00)
Constant	0.00***
	(0.00)
N	1350
R^2	0.232

Notes: ΔVote 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

with the largest increase in vote share. Nevertheless, since both shifts in similarity and electoral outcomes vary across constituencies, we interpret the result as a causal effect of terrorist attacks on changes in vote share. One, however, should refrain from giving a numerical interpretation to the magnitude of the coefficient: the reasons are the possible sources of bias which suggests that any effect we measure could underestimate the true effect of terrorist incidents on the support for radical parties.

Heterogeneity in Similarity Shifts

As explained, although our events are exogenous to any local characteristics, one interesting question is what local characteristics amplify or dampen the reaction of language shifts to terrorist attacks. Identifying the right set of independent variables which could possibly be correlated with this effect is not obvious. In the future we want to further analyze this issue using more sophisticated machine learning (really, specification search) techniques. So far 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. According 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 socioeconomic and demographic conditions in a typical German electoral constituencies. Note that this set might not be the optimal set for all parties. At this point, however, we simply investigate a correlation that can provide evidence on whether our measured effect is closely related to standard variables considered in studies of electoral outcomes.

Results are presented in Table 5 and can be easily summarized. As a dependent variable in each column we considered the cumulated shifts in similarity towards a party $\Pi_i^p = \{\pi_{3i}^{11p}\}$. None of the considered variables (share of people above 60 of age, share of foreign population, disposable income per inhabitant, amount of craftsmen firms, unemployment rate in 2017, share of high school graduates from the highest track in the German education system, share of employees in manufacturing and an indicator of East German constituencies) seems to be correlated with the size of reaction of terrorist attacks.

Robustness

Placebo Tests

Are the events we have selected really special, or are they similar to other event not related to terror or culture? In other words, is our model capturing strong points of discontinuities or would any other, non-political event produce the same result? We estimate the same discontinuous growth model to compute shifts in similarities for events which we argue are orthogonal to any political

Table 5: Investigating Heterogeneity in Similarity Shifts

	AfD	CDU	CSU	Die Grünen	Die Linke	FDP	SPD
Age (60+)	0.00	-0.00	0.00	-0.00	-0.00	-0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Foreign Population (%)	-0.00	0.00	0.00	-0.00	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Disp. Income (pc)	0.00	0.00	-0.00	0.00	0.00	0.00	0.00*
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Craftsmen Firms	0.00	-0.00	-0.00	0.00	0.00	-0.00	-0.00*
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Unemployment 2017 (%)	0.00	0.00	-0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
High Education (%)	0.00	0.00	0.00	0.00	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Manufacturing (%)	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
East	-0.00	0.00	0.00	-0.00	-0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N	225	189	36	225	225	225	225
R^2	0.023	0.082	0.086	0.028	0.028	0.026	0.038

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

context. We use sport events, namely the domestic German Cup finals (*DFB Pokalfinale*) in 2016 and 2017, as well as two matches during the UEFA Euro Cup in 2016. The results for shifts in similarity towards the AfD of these "placebo analysis" are presented in Figure 5. The analyses does not show a comparable increase in similarity towards the AfD relative to the events examined in the former section. For the fourth placebo event (*DFB Pokal Finale*, May 27, 2017), we do observe a discontinuity, but this event occurred just five days the Manchester bombing. Hence, we should not be surprised observing at least a small discontinuity.

Changes in Party Language

One natural concern with our method 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 one agent, the party account, and the public. In our analysis we treat essentially 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. However, we doubt that this is a relevant factor in our context. First, we exploit variation in language happening the exact day of unexpected attacks. If party activism, rather than voter attitudes, was explaining our findings, this would mean that the AfD account would be collecting reactions of Twitter users in all or most of German constituencies in real time and adapt also in real time Tweet content to the corpus of Tweets in each constituency. This hypothesis, applied to all events we consider, seems implausible.

Another possible threat would come if both parties and Twitter users tweeted about the details of events in real time, thus generating a convergence in language used. This however would not
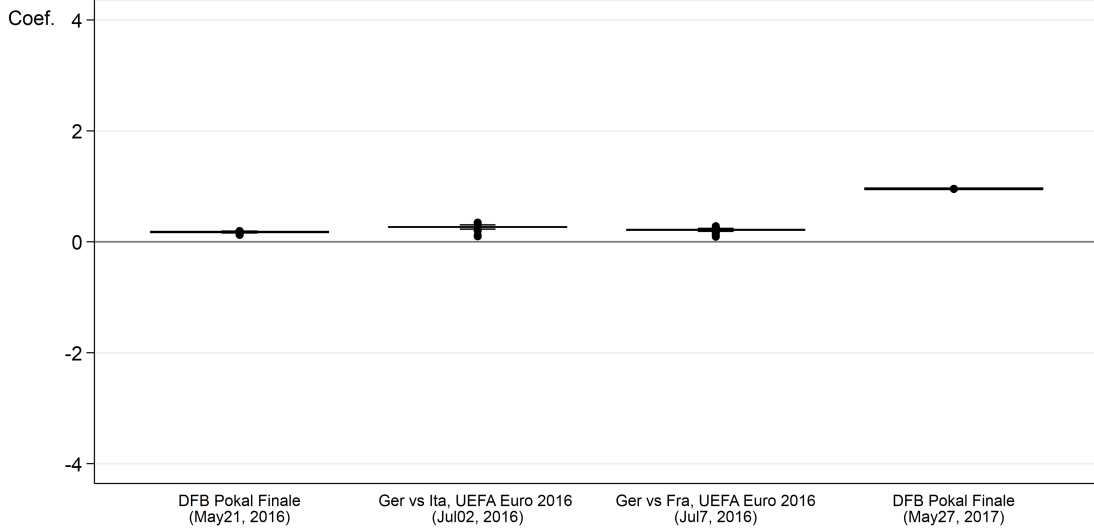
Figure 5: Placebo Events

explain why the large positive shifts in similarity are observed only for AfD and are reversed for other mainstream parties.

Finally, a possibility is that voters shifted their Twitter language towards that of AfD, but for reasons that are unrelated to attitudes about immigrants and perceived threats. For instance, if the views of AfD about the economy had become prevalent and widespread in society, rather than those on immigration and multiculturalism. We believe this is not the case. Starting from the right-wing turn of the party in July 2015, AfD communication on social media has been increasingly dominated by issues of nationalism and cultural identity. The relevance of cultural issues relatively to economic ones was increasing in the period we consider, as we now show. From the full set of daily Tweets sent from the AfD account, we counted the number of times that given words appeared. In Figure 6 we plot the share of daily Tweets in which words related to "Euro", "Greece", "Islam" and "Nation" were used. We follow Cantoni et al. (2017) in the choice of keywords. We plot 3 months moving averages of the daily share of Tweets for each keyword. Although the graph is not informative about the specific events, it shows a trend. The vertical lines mark the 2015 Congress which changed the identity of AfD and the Federal Elections of September 24, 2017. In the period in between, where all our events happen, the use of the word "Nation" remains constantly high, after having decreased before 2015, and increases markedly in 2017, while topics such as "Euro" and "Greece" decrease steadily over time. "Islam" is also increasingly used between 2015 and 2016, even if to a minor extent. We are thus confident that any change in language towards that of AfD on social media is likely to be driven by cultural topics.

9   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 generally xenophobic. The regularly 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 to disentangle economic and cultural factors behind populist vote.

In this paper we exploit the exogenous timing of large events of terrorism and crime to isolate

23

Figure 6: Frequency of Specific Words in AfD Tweets

their independent effects on peoples' alignment with the values promoted by a right-wing populist party. Using an algorithm based on followers networks to assign Twitter users to geographic constituencies and Natural Language Processing techniques, we show that unexpected terrorist attacks and large crime news move the language of peoples' tweets closer to that of the German right-wing party.

We argue that terrorist attacks and large-scale crimes committed by immigrants constitute information shocks which 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. The effect was to move 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 are an independent factor driving 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 and opposite shifts from left-wing SPD and centrist CDU, the two main parties at the time in government. We perform placebo experiments to show that these effects are not the result of generally salient events unrelated to politics. Finally, we show that predicted language shifts towards AfD and other parties following the events we analyze are not explained by standard demographic and economic variables.

At this stage of our research, we are cautious with respect to the interpretation of our findings. However, we emphasize that our evidence is overall consistent with the channel that we hypothesize. Specifically, shifts in language following the events we consider are unlikely to be driven by discourse over economic issues, which was being considerably downplayed in the social media messages of the AfD during the period studied. Our results are a first step towards understanding cultural backlash as a channel of transmission of radical right support, which has been traditionally difficult to study and which we can address in the appropriate way using new data and techniques.

References

Alesina, A., Miano, A., and Stantcheva, S. (2018), Immigration and Redistribution. *NBER Working Paper Series*.

Algan, Y., Guriev, S., Papaioannou, E., and Passari, E. (2017), The European Trust Crisis and the Rise of Populism. *Brookings Papers on Economic Activity.* 2017(2), 309–400.

An, J., Cha, M., Gummadi, K. P., Crowcroft, J., and Quercia, D. (2012), Visualizing Media Bias through Twitter.

Arzheimer, K. (2019), *The Eclectic, Erratic Bibliography on the Extreme Right in Western Europe.*

Autor, D., Dorn, D., Hanson, G., and Majlesi, K. (2016), Importing Political Polarization? *Working paper.*

Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., and Bonneau, R. (2015), Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological science.* 26(10), 1531–1542.

Barone, G., D'Ignazio, A., Blasio, G. de, and Naticchioni, P. (2016), Mr. Rossi, Mr. Hu and politics. The role of immigration in shaping natives' voting behavior. *Journal of Public Economics.* 136, 1–13.

Betz, H.-G. (1994), *Radical Right-Wing Populism in Western Europe.* New York: Palgrave Macmillan US and Palgrave Macmillan.

Bhatia, S., Lau, J. H., and Baldwin, T. (2016), Automatic Labelling of Topics with Neural Embeddings. *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 953–963.

Bliese, P. D. and Lang, J. W. B. (2016), Understanding Relative and Absolute Change in Discontinuous Growth Models. *Organizational Research Methods.* 19(4), 562–592.

Cantoni, D., Hagemeister, F., and Westcott, M. (2017), Persistence and Activation of Right-Wing Political Ideology. *Munich Discussion Paper.*

Card, D., Dustmann, C., and Preston, I. (2012), Immigration, Wages, and Compositional Amenities. *Journal of the European Economic Association.* 10(1), 78–119.

Colantone, I. and Stanig, P. (2018a), Global Competition and Brexit. *American Political Science Review.* 112(02), 201–218.

— (2018b), Replication Data for: The Trade Origins of Economic Nationalism: Import Competition and Voting Behavior in Western Europe. *American Journal of Political Science.* (62), 936–953.

— (2018c), The trade origins of economic nationalism: Import competition and voting behavior in Western Europe. *American Journal of Political Science.* 62(4), 936–953.

Curini, L., Iacus, S., and Canova, L. (2015), Measuring Idiosyncratic Happiness Through the Analysis of Twitter: An Application to the Italian Case. *Social Indicators Research.* 121(2), 525–542.

Dal Bó, E., Finan, F., Folke, O., Persson, T., and Rickne, J. (2018), Economic Losers and Political Winners: Sweden's Radical Right.

Der Bundeswahlleiter (2017), *Wahlkreise und zugeordnete Gemeinden bei der Wahl zum 19. Deutschen Bundestag.* Wiesbaden.

Dustmann, C., Eichengreen, B., Otten, S., Sapir, A., Tabellini, G. E., and Zoega, G. (2017), *Europe's trust deficit: Causes and remedies.* Vol. 1. Monitoring international integration. London: CEPR Press.

Dustmann, C., Vasiljeva, K., and Damm, A. P. (forthcoming), Refugee migration and electoral outcomes. *Review of Economic Studies.*

European Parliament and the Council (2003), *Regulation (EC) No 1059/2003 of the European Parliament and of the Council of 26 May 2003 on the establishment of a common classification of territorial units for statistics (NUTS).*

Foreign Policy (22.10.2015), Nothing Can Take Down Angela Merkel — Except 800,000 Refugees.

Franz, C., Fratzscher, M., and Kritikos, A. S. (2018), *German right-wing party AfD finds more support in rural areas with aging populations.*

Frees, B. and Koch, W. (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.

Funke, M., Schularick, M., and Trebesch, C. (2015), Going to Extremes: Politics after Financial Crisis, 1870-2014.

Gentzkow, M. and Shapiro, J. M. (2010), What Drives Media Slant? Evidence From U.S. Daily Newspapers. *Econometrica.* 78(1), 35–71.

Giani, M. (2017), Don't Ban the Muslims: Public Reaction to Terror Attacks.

Goerres, A., Spies, D. C., and Kumlin, S. (2018), The Electoral Supporter Base of the Alternative for Germany. *Swiss Political Science Review.* (24), 246–269.

Golder, M. (2003), Explaining Variation In The Success Of Extreme Right Parties In Western Europe. *Comparative Political Studies.* 36(4), 432–466.

Gorodnichenko, Y., Pham, T., and Talavera, O. (2018), *Social Media, Sentiment and Public Opinions: Evidence from #Brexit and #USElection.* Cambridge, MA: National Bureau of Economic Research.

Halberstam, Y. and Knight, B. (2016), Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter. *Journal of Public Economics.* 143, 73–88.

Halla, M., Wagner, A. F., and Zweimüller, J. (2017), Immigration and voting for the far right. *Journal of the European Economic Association.* 15(6), 1341–1385.

Hangartner, D., Dinas, E., Marbach, M., Matakos, K., and Xefteris, D. (2017), Does Exposure to the Refugee Crisis Make Natives More Hostile? *American Political Science Review*, 1–14.

Hansen, M. A. and Olsen, J. (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.

Huberty, M. (2015), Can we vote with our tweet? On the perennial difficulty of election forecasting with social media. *International Journal of Forecasting.* 31(3), 992–1007.

infratest dimap (2018), *Umfragen & Analysen: Bundesländer.*

Inglehart, R. F. and Norris, P. (2016), Trump, Brexit, and the Rise of Populism: Economic Have-Nots and Cultural Backlash. *HKS Faculty Research Working Paper Series.*

Jann, O. and Schottmüller, C. (2018), Why Echo Chambers are Useful. *University of Oxford, Department of Economics, Discussion Paper Series.*

Jin, Z., Juan, C., Han, G., Yongdong, Z., Yu, W., and Jiebo, L. (2017), Detection and Analysis of 2016 US Presidential Election Related Rumors on Twitter. *Social, Cultural, and Behavioral Modeling.* Ed. by D. Lee, Y.-R. Lin, N. Osgood, and R. Thomson. Cham: Springer International Publishing, pp. 14–24.

Kitschelt, H. (1995), *The radical right in Western Europe: A comparative analysis.* Ann Arbor: Univ. of Michigan Press.

Knigge, P. (1998), The ecological correlates of right-wing extremism in Western Europe. *European Journal of Political Research.* 34(2), 249–279.

Kobold, K. and Schmiedel, S. (2018), Wahlverhalten bei der Bundestagswahl 2017 nach Geschlecht und Alter: Ergebnisse der repräsentativen Wahlstatistik. *WISTA - Wirtschaft und Statistik.* (3), 142–156.

Le, Q. and Mikolov, T. (2014), Distributed representations of sentences and documents. *International conference on machine learning*, pp. 1188–1196.

Legewie, J. (2013), Terrorist Events and Attitudes toward Immigrants: A Natural Experiment. *American Journal of Sociology.* 118(5), 1199–1245.

Liu, B. (2012), Sentiment Analysis and Opinion Mining. *Synthesis Lectures on Human Language Technologies.* 5(1), 1–167.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013), Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781.*

National Consortium for the Study of Terrorism and Responses to Terrorism (2018), *Global Terrorism Database.*

Obschonka, M. and Fisch, C. (2018), Entrepreneurial personalities in political leadership. *Small Business Economics.* 50(4), 851–869.

Oesch, D. (2008), Explaining Workers' Support for Right-Wing Populist Parties in Western Europe: Evidence from Austria, Belgium, France, Norway, and Switzerland. *International Political Science Review.* 29(3), 349–373.

Park, S., Lee, J., Ryu, S., and Hahn, K. S. (2015), The Network of Celebrity Politics. *The ANNALS of the American Academy of Political and Social Science.* 659(1), 246–258.

Politico (20.12.2016), Berlin attack exposes Angela Merkel's right flank.

Ravi, K. and Ravi, V. (2015), A survey on opinion mining and sentiment analysis: Tasks, approaches and applications. *Knowledge-Based Systems.* 89, 14–46.

Reuters (10.01.2016), Growing scale of Cologne attacks stokes German debate on migrants.

Rodrik, D. (2018), Populism and the Economics of Globalization. *Journal of International Business Policy.* 1(1-2), 12–33.

Roth, A. and Wolff, G. B. (5.10.2017), *What has driven the votes for Germany's right-wing Alternative für Deutschland?*

Safiullah, M., Pathak, P., Singh, S., and Anshul, A. (2017), Social media as an upcoming tool for political marketing effectiveness. *Asia Pacific Management Review.* 22(1), 10–15.

Spiegel Online (23.09.2013), German Euroskeptics' Breakthrough Moment: Next Stop, Brussels?

Steinmayr, A. (2016), Exposure to Refugees and Voting for the Far-Right: (Unexpected) Results from Austria. *IZA Discussion Paper.*

Sterling, J. and Jost, J. T. (2018), Moral discourse in the Twitterverse: Effects of ideology and political sophistication on language use among U.S. citizens and members of Congress. *Journal of Language and Politics.* 17(2), 195–221.

Süddeutsche Zeitung (10.07.2016), 1200 Frauen wurden Opfer von Silvester-Gewalt.

Vaccari, C., Valeriani, A., Barberá, P., Bonneau, R., Jost, J. T., Nagler, J., and Tucker, J. (2013), Social media and political communication. A survey of Twitter users during the 2013 Italian general election. *Rivista italiana di scienza politica.*

Vaccari, C., Valeriani, A., Barberá, P., Jost, J. T., Nagler, J., and Tucker, J. A. (2016), Of Echo Chambers and Contrarian Clubs: Exposure to Political Disagreement Among German and Italian Users of Twitter. *Social Media + Society.* 2(3), 205630511666422.

Vertier, P. and Viskanic, M. (2018), Dismantling the "Jungle": Migrant Relocation and Extreme Voting in France. *CESifo Working Papers.*

Zeit Online (5.01.2016), Polizei gibt Fehler in Silvesternacht zu.

Appendix A

Measured Similarity

As explained before, inference based on changes in measured similarity at the day of events are subject to threats of validity (compare Section 7). In short, assuming that the people using Twitter during the immediate influence of a terrorist attack represent the overall (Twitter) population of a given constituency appears to be not realistic. Nevertheless, for the sake of completeness Figure 7 shows the measured changes in similarity between all constituencies and the AfD around the considered events. Notice that we have a high variation across constituencies, showing both positive and negative changes.

The median change, indicated by the horizontal line within each boxplot, appears to be pretty close to zero. Again, this change reflects the immediate influence of the events. It is computed by taking the difference between the value of similarity at the day of an event and the day before. However, since attacks usually happened at night, and it may have taken some hours until a considerable amount of news was spread, results are likely to look different on other days after the events.

Possible arbitrariness in the choice of a reference point after each single event to compute this difference raises incentives to use a more sophisticated strategy, which is exactly what is achieved by making use of the discontinuous growth model.

Discontinuous Growth Model

Variables considered so far:

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*: The second change variable (Event) is coded 0 prior to the event and 1 after the event. The parameter estimate associated with this variable ($\pi^p_{3ki} = \pi^p_{3k} + r^p_{3ki}$) determines the degree to which the intercept was altered after the event.


Figure 7: Changes in Measured Language Similarity to the AfD by Constituency

Table 6: Coding of Time Variables - Multiple Events

$Time$	$Time^2$	E_1	E_2	$Reset_1$	$Reset_1^2$	$Reset_2$	$Reset_2^2$
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
...	...	...	...	...	...	...	...

Note that here we have the first event occurring in period 5, the second event occurring in period 10.

4. *Reset*: The third change variable is coded 0 at the period of which the event first occurs and increases with each subsequent period. It functions like a new starting time trend. The parameter estimate $\pi_{4ki}^p = \pi_{4k}^p + r_{4ki}^p$ indicates the degree to which the event alters the slope ($\pi_{1i}^p = \pi_1^p + r_{1i}^p$) for the time variable.

5. *Reset²*: Similar to before with a quadratic change variable.

6. *Year*: Standard year fixed effect being equal to 1 in 2016 and 0 elsewhere. Only one year included due to high multicollinearity.

For analyzing multiple events, we simply introduce multiple variables for events and changes. Consider the following Table 6 as overview on the variables.

In our results we estimate the discontinuous growth model introduced in section 7 using maximum likelihood with a linear and quadratic time trend, the set of events as changes of intercept and in linear and quadratic change in slope. We include random coefficients (r) for the linear component of time trend and the shifts in similarity. The estimated equation for measured similarity $simil_{i,t}^p$, for a given party p in constituency i at time t is in this case:

$$
\begin{aligned}
simil_{ti}^p = {}& \pi_{0i}^p + \pi_1^p Time_t + r_{1i}^p Time_t + \pi_2^p Time_t^2 + r_{2i}^p Time_t^2 \\
& + \sum_{k=1}^{11} \pi_{3k}^p E_{kt} + \sum_{k=1}^{11} r_{3ki}^p E_{kt} \\
& + \sum_{k=1}^{11} \pi_{4k}^p Reset_{kt} + \sum_{k=1}^{11} r_{4ki}^p Reset_{kt} \\
& + \sum_{k=1}^{11} \pi_{5k}^p Reset_{kt}^2 \\
& + \pi_6 Year + \epsilon_{ti}^p
\end{aligned}
\tag{4}
$$