Online Appendix

A Party Political Communications

In this section we discuss the scraping of party political Press releases of both the mainstream parties (CDU/CSU, SPD, Greens, FDP) and the AfD (and the Left), and the retrieval of the refugee-related content from those sources.We used the press releases of the parliamentary groups (Fraktionen) for the party political communication for two reasons. First, the parties as such did not publish many press releases independent from the parliamentary groups, so there is only very few data to extract from the party press releases. Secondly, the parliamentary groups need to react to recent political developments quickly and often use press releases as their preferred mode, making them comparable to newspaper articles (which often need to rely on press releases by the parties). The AfD, however, used press releases very often as tool of party communication before 2017. Thus, we included their party press releases in our analysis. As a robustness check, we re-ran the analysis for the time span after 2017 with the press releases of the parliamentary group.

In order to scrape the press releases from the political parties, we identified the homepages where the press releases were stored in a first step. For the CDU/CSU, SPD, and the Left, this was a very easy task, since they stored all press releases from the last years openly on their homepages. As such, they could be easily scraped using simple web-scraping algorithms. For the FDP, we had to consider that they were not represented in the Bundestag in 2013-2017. Here, we decided to refrain from using party press releases - and excluded the FDP from 2013-2017. Furthermore, the FDP hid their press releases after 2017 in a JSON container, which had to be identified manually. From this, we could apply the same web scraping tools than for the other parties. The Greens were quite complicated to scrape, since they both hid their press releases in JSON container and removed all press releases earlier than 2017 from their homepage. When checking the web.archive, we recognized that the old press releases have not been archived by the web.archive as well. Luckily, we could access the data from a different project in 2018. For the AfD, we also scraped the press releases from their homepage. However, we faced a problem while scraping, as the party has deleted all press releases earlier than January 2015 from their homepage. In order to fill the coverage we retrieved the remaining press releases from the web-archive. Table 3 shows the number of press releases we obtained for each political party.

Next, in order want to capture the meaning of each text, we employ indomain word embeddings that have been trained on German political parties manifestos from the Manifesto Project. For each article, we average the embeddings of all the words it contains, getting an overall document embedding; we process our query the same way. To rank each document given the query, we measure their semantic similarity as the cosine similarity of the two vectors (between the query and each press release). We choose cutoffs between 0.35

Table 3

Political Party	Number of Scraped Press Releases	Number of Refugee-Relevant Releases
CDU/CSU	4,627	330
SPD	3,800	184
AfD	2,007	461
Greens	4,345	226
The Left	5,700	496
FDP	1,711	64
Total	21,498	1,642

and 0.39 of the similarity score, relying on several evaluation processes, where we took random samples of releases that fell between different cutoffs points to check how many articles that were relevant to refugee narratives we have captured or missed. To extend our results, we additionally add any release that mentions at least one of our refugee-related keywords in the title of the release. Table 3 shows the numbers of our refugee related press releases for each party.

Other than the srcaping of party communications, which was done in R, the remaining analysis was done in Python, with the help of the NTLK, Spacy and Gensim packages.

We conducted an examination of the quality of the corpus. In order to evaluate recall, we selected at random 500 party communications. Human coders uncovered 44 refugee-relevant documents. Of these, 28 had been identified by our automated search, yielding a recall rate of 64 %. Next, we selected a random sample of 150 automatically-identified party communications. We found that 90 % of them are actually refugee-relevant, and this is our precision.

A.1 Keywords for Refugee Topic

Flüchtling, Flüchtlings, Fluechtlings, fluechtlings, fluechtlinge, Fluechtling, Flüchtlinge, Flüchtlingen, Fluechtlingen, fluechtlinge, fluechtlingen, Flüchtlingsandrang, Fluechtlingsandrang, Flüchtlingsstatus, Flüchtlingsunterkunft, Flüchtlingsströme, Flüchtlingsströmen, Fluechtlingsstroeme, Fluechtlingsstroemen, Fluechtlingsunterkunft, Flüchtlingsunterkunfte, Fluechtlingsunterkunfte, Fluechtlingsstatus Fluechtlinge, Fluechtlingspolitik, Flüchtlingspolitik, Flüchtlingskrise, Fluechtlingskrise, Fluechtlingskinder, Flüchtlingskinder, Flüchtlingskriminalität, Antimigrationspolitik, Anti-migrationspolitik, Fluechtlingskriminalitaet, Fluchtgesetze, Flüchtlingsgesetze, Fluechtlingsgesetze, Asyl-Gesetzgebung, Asylgesetzgebung, Asyl-Gesetzgebung, Asylgesetze, Asylanträge, Asylsuchende, Asylantraege, Asylpaket, Asyl, Asylpolitik, Asylbewerber, Asylbewerbern, Migranten, Migrant, Migrantenpolitik, Migrationskrise, migrationskrise, Einwanderer, Einwanderung, Flucht, Geflüchteten, Geflüchtete, Gefluechteten, Gefluechtete, Flüchtlingszustrom, Fluechtlingszustrom, Migrationsproblem, Migrationsprobleme, Migration, Migrationsproblemen, Migrationsproblematik, Kriegsflüchtlingen, Kriegsfluechtlingen, Kriegsflüchtlinge, Kriegsfluechtlinge, Flüchtlingsstrom, Fluechtlingsstrom, Flüchtlings-Obergrenze, Fluechtlings-Obergrenze, Flüchtlingsboot, Fluechtlingsboot, Fluchtursachenbekampfung, Fluchtursachen

A.2 Sentiment Analysis

Here we include sentiment score by source, for all sources, and information on validation. We performed an evaluation of the sentiment analysis by selecting 25 documents from each source and having human coder assign -1,0,1 for negative, neutral and positive sentiments toward refugees in document. We used the data to generate graphs for how different outlets fare on the human coding (STATA), and for what the sentiment scores look like for different values of the human coding (R).

Figure 6: Sentiment Analysis: AfD vs Mainstream parties on Refugees. Proportion positive words (positive over positive+negative over month). Based on 1265 total number of documents.

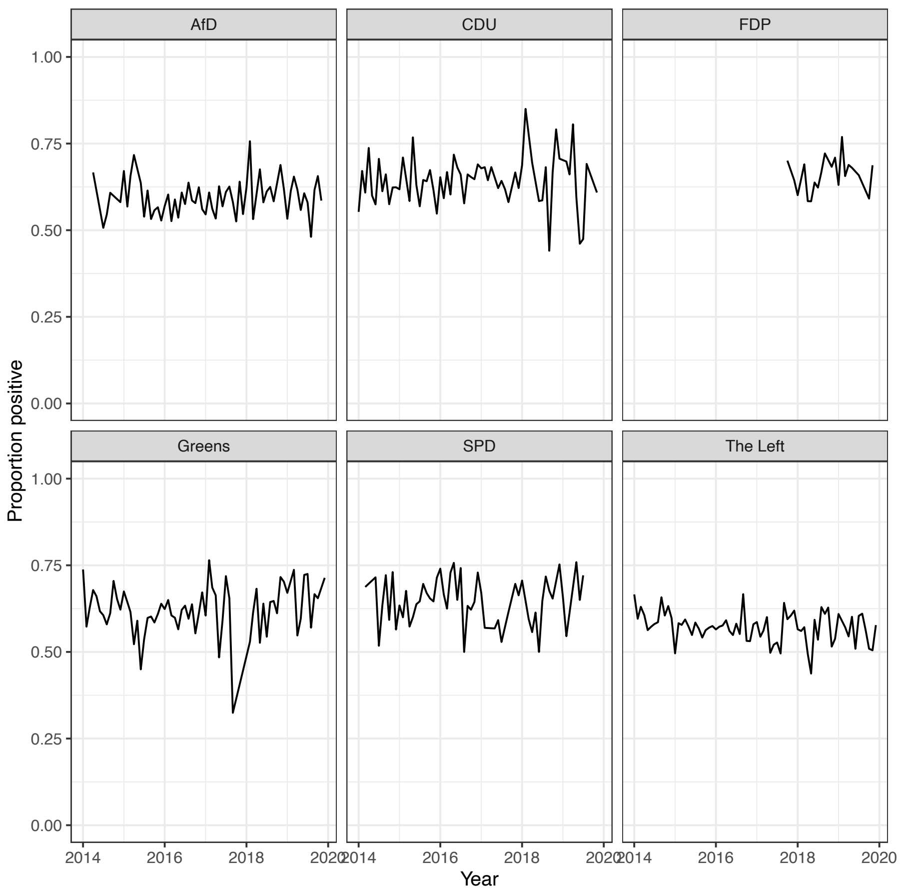

B Mass Media Data

We rely on archived version of select newspaper websites to collect our newspaper data, preserved by the Internet Archive (IA), a nonprofit digital library that since 1996 has been consistently capturing and making available for future studies as much of the public web as possible (Kahle 1997). As of September 2018, IA has collected more than 25 petabytes of web data. Through its API, we collect preserved articles for Frankfurter Allgemeine Zeitung (FAZ), Süddeutsche Zeitung (SZ), Welt, Tageszeitung (TAZ), Bild, RT Deutsch, and Sputnik.de, in the interval between January 1st, 2014 and December 31st, 2018 (note again that RT Deutsch and Sputnik.de come into operation on November 15, 2014.) To retrieve the refugee-related articles from the corpora of newspaper we have collected, we follow the same procedure that was explained in above in the appendix for filtering the refugee-related press releases.

We retrieve online published articles from the Internet Archive of the following newspaper: Bild, FAZ (Frankfurter Allgemeine), SZ (Süddeutsche Zeitung), Welt, TAZ (Tageszeitung), RT.de, and Sputnik.de (de.sputniknews), that were featured on their homepages between the years 2014 and 2018. We scrape all daily snapshots of each newspaper's homepage from the IA, extract all links to the presented articles and collect each of these stories.

We use the same procedure described in about party political communication above to capture the refugee-relevant content from our collection. We use the same dictionary/query of relevant terms that refugee related. We also use the same text-preprocessing technique for the query and each of the articles: we tokenize in single words, remove stopwords, and punctuation and numbers.

To move beyond simple lexical overlap between terms in the query and in the articles and in order to capture the meaning of each text, we employ the indomain word embeddings. We then measure the cosine similarity between each of the articles and the query. We choose a cutoff of 0.4 similarity score.[41] Here also we relied on several evaluation processes to find the correct cut-off point. To extend our results, we additionally add any article that mentions at least one of our refugee-related keywords in the url or in its tags of the main topics of the article.

To further examine the quality of our results, we conducted an additional evaluation that compared randomly-selected refugee-related articles that were retrieved manually from the web and others that were captured through our computational algorithm. Through this evaluation, we determine that our collection represents comprehensive coverage of refugee relevant news, which is representative of the narratives the newspapers published at a particular point of time. Our evaluation of the corpus shows that its quality is high: we miss

[41]FAZ 0.35

virtually no stories, a recall approaching 100%. The precision, depending on the outlet, lies between 80 to 98 per cent of relevant stories to the refugee-related topic.

The numbers of the refugee related articles in comparison to the total scraped articles for each outlet is reported in Table 4.

Table 4

Newspaper Source	Number of Scraped Articles	Number of Refugee-Relevant Articles
Bild	69,931	917
RT Deutsch	40,216	745
Sputnik.de	148,187	3,324
FAZ	123,200	9,199
SZ	128,396	3,832
TAZ	68,154	3,361
Welt	150,394	4,073
Total	728,478	25,450

B.1 Sentiment Analysis

Figure 7: Sentiment Analysis: German (FAZ, TAZ, Welt, Bild, SZ) and Russian Media (Sputnik and RT) on Refugees. Proportion positive words (positive over positive+negative over month). Based on 25,450 total number of documents.

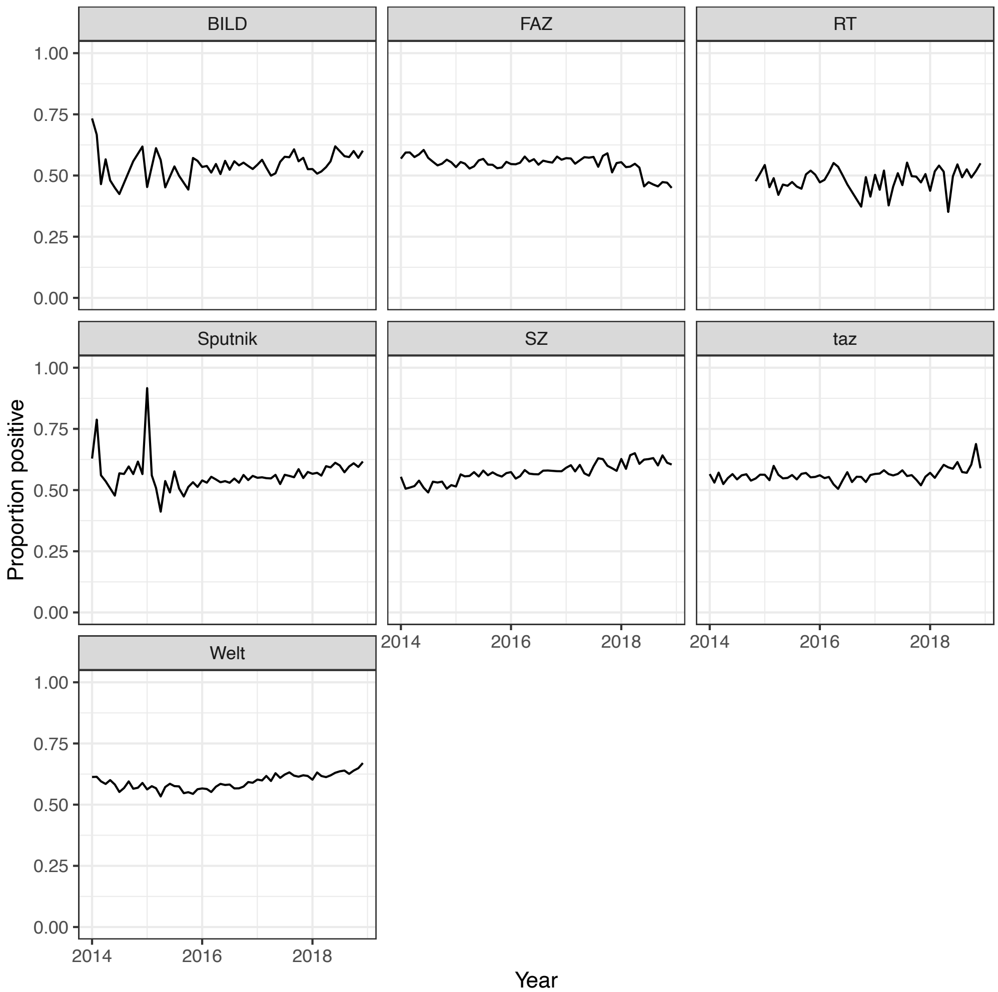

C Validation of Sentiment Scores

Figure 8: Validation of sentiment, random sample of 275 texts

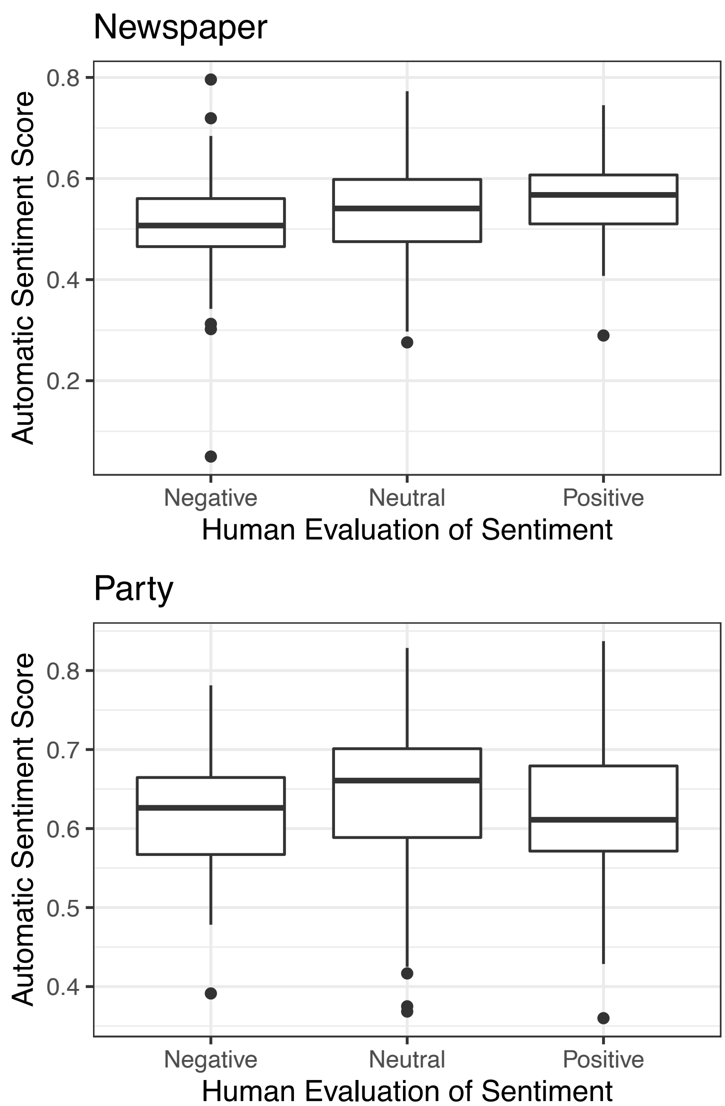

Figure 9: Validation of sentiment by outlet, random sample of 275 texts

(a) Party Communications

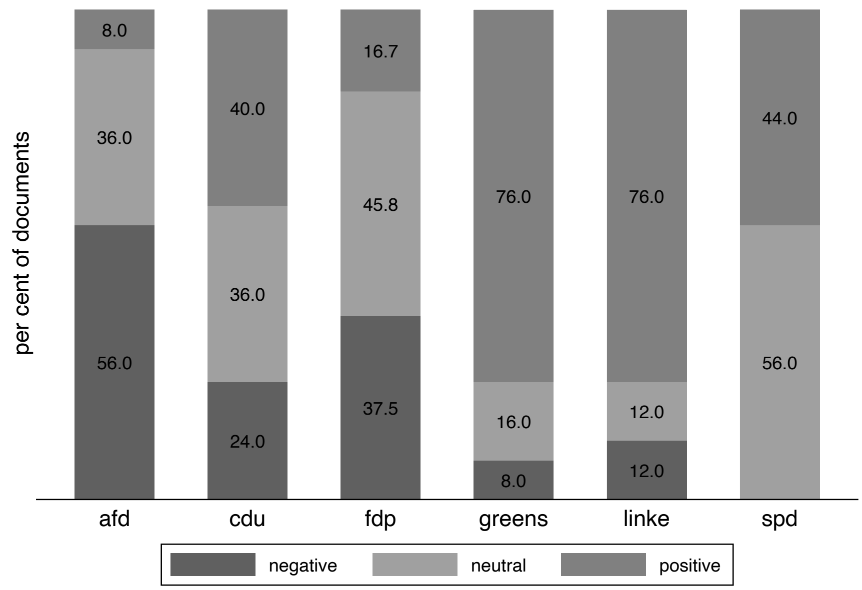

(b) Mass Media

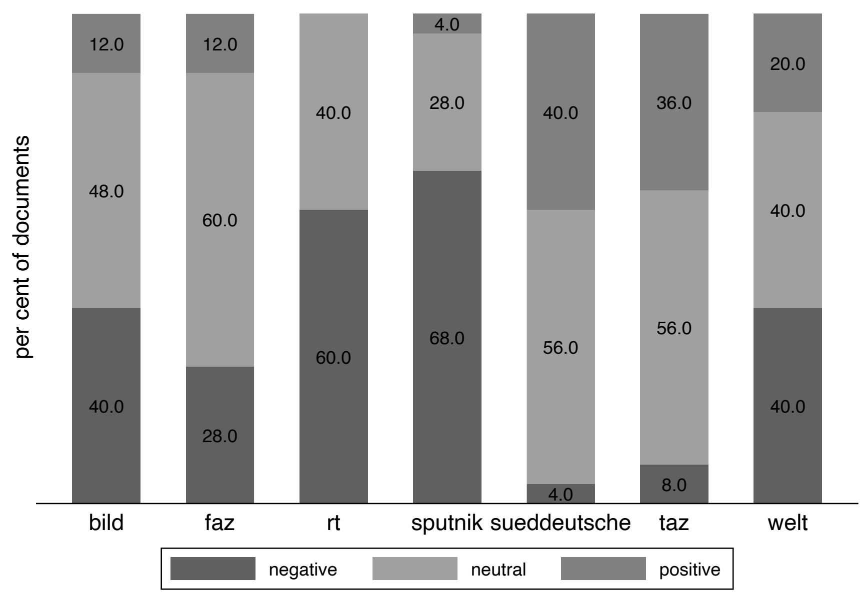

D Procedure and Keywords for Similarity to Conspiratorial Language

We generate a set of keywords for the conspiracy topic based on our reading of a broad array of articles across outlets (see below the keywords). We adopt a semantic text similarity approach (Šarić et al. 2012)), relying on word embeddings. We compute a similarity score between each article's words and the topic words using our indomain embeddings. For comparison purposes we also use wikipedia embeddings that have been trained by Facebook researchers on Wikipedia content to compute the similarity score, mainly to see if results hold using different embeddings. This described technique, is the same as what we adopted to identify a story as refugees-related, except now we seek to sub-classify the refugee-relevant stories by computing their relevance to topic seed-words. We interpret the semantic similarity score between a story and the topic words (with a range of -1 to 1) as evidence of more emphasis on conspiracy. We rely on this approach instead of other text analysis techniques (e.g. classification, topic modeling, or sentiment analysis) as it allows us of a) defining in advance and having full control over the topics of interest through their lists of keywords (as opposed than when working with topic models), b) avoiding expensive manual coding of training examples (necessary for text classification) and c) modeling and identifying relevance to a specific topic instead of sentiment towards it (which is problematic to capture).

We believe that in-domain embeddings are better-suited for texts with political communication. Wikipedia embeddings exist for more words.

Keywords:
geopolitisch, Einflussnahme, Marionette, Weltlenker (**), Verschwörung (*), Geheimdokumente (*), verheimlichen, Meinungstyrannen (**), Soros (*), Euromaidan (*), Orange Revolution, Open Society

Note: words marked with * do not have embeddings in the In-Domain corpus (Party Manifestos) and words marked with ** have embeddings neither in In-Domain nor in the Wikipedia/Facebook corpus

Figure 10 shows a comparison of the two types of embeddings. Russian sources features significantly higher mean conspiracy score in both approaches.

To further validate and illustrate the similarity scores approach, we took ten randomly selected stories from the bottom quartile and the top quartile for the topic, from each outlet we study. We then manually check whether the story covers the topic, as defined above. Table 5 shows the per cent stories that fit as we move from the bottom to the top quartile on the relevant score. The results are encouraging. For conspiracy, stories that are rarer, we see an increase from 0 per cent, to up to 30 per cent. Thus, higher similarity scores do pick out stories that are relevant to that topic.

Figure 10: Conspiracy language: comparing embeddings

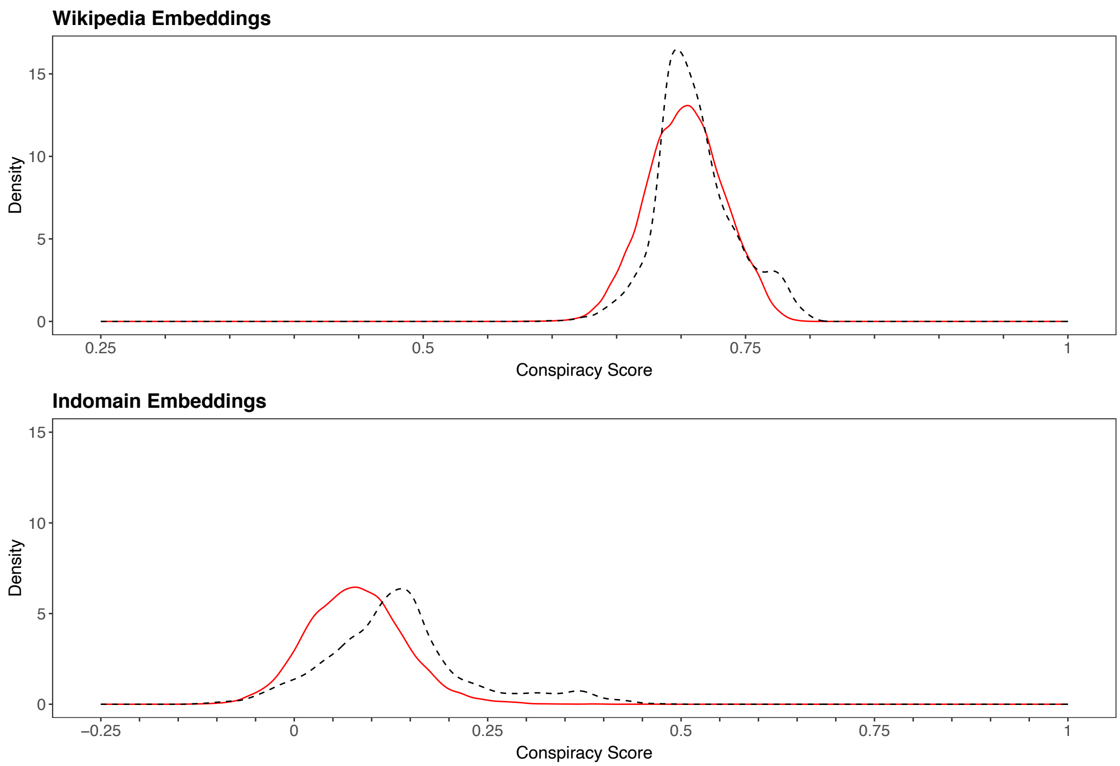

Table 5: Effect of Moving from Bottom to Top Quartile of Similarity Score on Percent Relevant Articles Discovered in Outlet

	(1) Bild	(2) FAZ	(3) SZ	(4) Welt	(5) TAZ	(6) Sputnik	(7) RT
Conspiracy	0.02 0%	0.05 0%	0.05 0%	0.02 0%	0.01 0%	0.07 0%	0.17 0%
	0.10 10%	0.13 30%	0.12 0%	0.11 0%	0.07 0%	0.26 10%	0.32 20%

Note: Upper/lower limits of lower/upper quartile listed to the left of % relevant; based on random sample of 420 stories

Figure 11: How does dropping a word affect conspiracy scores? Ablation analysis

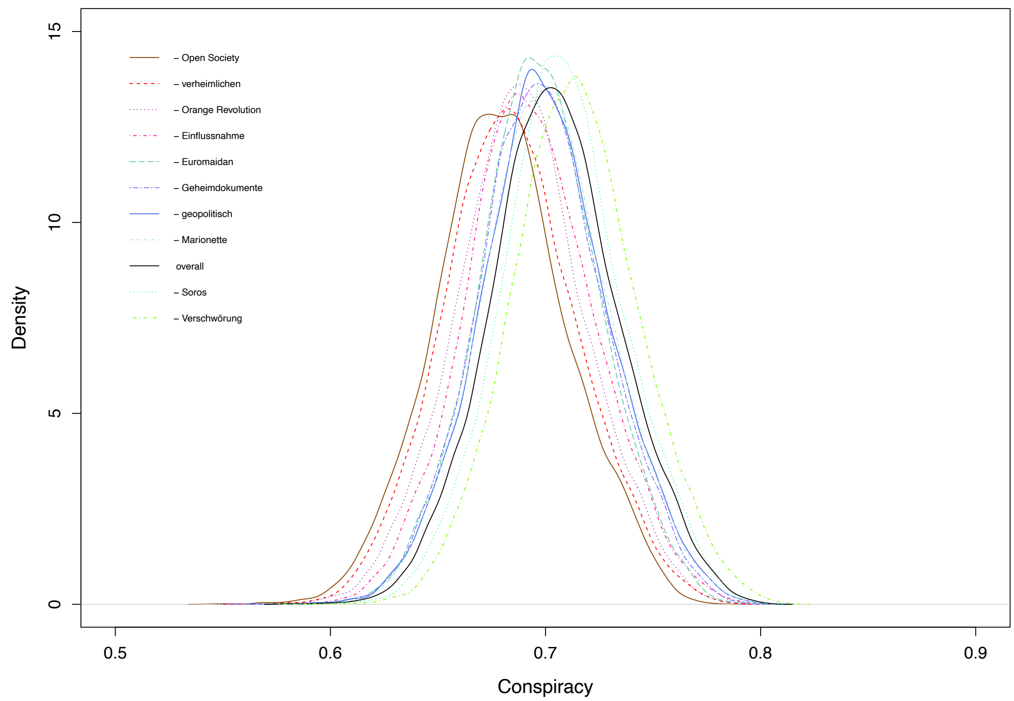

E Event analysis

Figure 12 and Table 6 show results from event-study and a simple t-test from a longer window of 6 weeks before and after the election as the event window (and so -12 to -6 and 6 to 12 being the non-events). As expected, longer windows introduce noise and weaken the results somewhat. In the event-study, the RT effect barely misses statistical significance for the federal election. It does retain significance at the 95 % level in the simple difference-of-means t-test approach. Interestingly, Sputnik appars to play a stronger role in boosting coverage in two consequential Land elections.

Figure 12: Event study: deviation of observed from expected daily refugee stories, 95 % CI. DE=Federal Election.

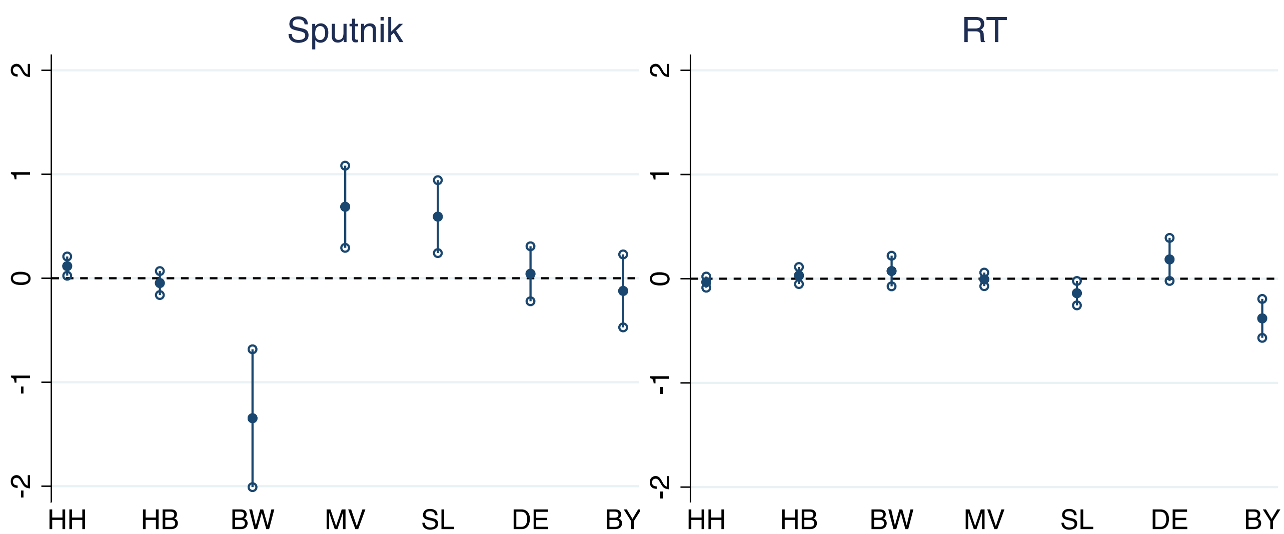

Table 6: Results of difference of means t-test on number of daily stories on refugees close to elections (event=1) and further out (event=0), by election/event.

source	event	days event=0	daily stories event=0	days event=1	daily stories event=1	two-sided p-value
RT Deutsch	HH	84	0.13	83	0.108	0.673
RT Deutsch	HB	84	0.178	83	0.204	0.708
RT Deutsch	BW	84	0.25	83	0.421	0.113
RT Deutsch	MV	84	0.214	97	0.144	0.269
RT Deutsch	SL	84	0.392	132	0.371	0.816
RT Deutsch	DE	84	0.345	104	0.653	0.042
RT Deutsch	BY	65	1.446	96	1.083	0.04
Sputnik.de	HH	84	0.3	83	0.156	0.117
Sputnik.de	HB	84	0.285	83	0.253	0.7508
Sputnik.de	BW	84	4.428	83	4.963	0.496
Sputnik.de	MV	84	1.535	97	2.402	0.005
Sputnik.de	SL	84	1.857	132	1.969	0.697
Sputnik.de	DE	84	1.583	104	1.701	0.626
Sputnik.de	BY	65	2.369	96	1.906	0.177