UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

Dipartimento di Economia Marco Biagi

DEMB Working Paper Series

N. 146

Does fake news affect voting behaviour?

Michele Cantarella*, Nicolò Fraccaroli**, Roberto Volpe***

June 2019

* Fondazione Marco Biagi, University of Modena and Reggio Emilia and European Central Bank
E-mail: michele.cantarella@unimore.it
** University of Rome Tor Vergata
*** Italian Ministry of Economic Development (MISE)

ISSN: 2281-440X online

Does fake news affect voting behaviour?

Michele Cantarella*,[1], Nicolò Fraccaroli[2], and Roberto Volpe[3]

[1]Fondazione Marco Biagi, University of Modena and Reggio Emilia and European Central Bank

[2]University of Rome Tor Vergata

[3]Italian Ministry of Economic Development (MISE)

6th June 2019

Abstract

Over the last decade, the erosion of trust in public institutions and traditional media sources have been proceeding in parallel. As recent developments in media consumption have led to a proliferation of politically charged online misinformation, it is no wonder that many have been questioning whether the spread of fake news has affected the results of recent elections, contributing to the growth of populist party platforms. In this work, we aim to quantify this impact by focusing on the causal effect of the spread of misinformation over electoral outcomes in the 2018 Italian General elections. We exploit the presence of Italian and German linguistic groups in the Trento and Bolzano/Bozen autonomous provinces as an exogenous source of variation, assigning individuals into distinct filter bubbles each differently exposed to misinformation. To do so, we introduce a novel index based on text mining techniques to measure populism. Using this approach, we analyse the social media content of each party and their leaders over the course of the electoral campaign for the 2013 and 2018 elections. We then collect electoral and socio-demographic data from the region and, after constructing a proxy for exposition to misinformation, we measure the change in populist vote across the two groups in-between the two general elections, using a combination of difference-in-difference and two-stage-least-squares inference methods. Our results indicate that misinformation had a negligible and non-significant effect on populist vote in Trentino and South Tyrol during the Italian 2018 general elections.

Keywords: *Fake News, Political Economy, Electoral Outcomes, Populism*

*Corresponding author. Email: michele.cantarella@unimore.it.

1 Introduction

Over the last several years, in parallel with the expansion of online communities outside of their previously younger demographics, fake news has proliferated online. As recent elections across the globe have been characterised by a swing towards populism – from the 2016 US presidential elections to the 2018 Italian parliamentary elections – it is no surprise that the spread of fake news has been held responsible for affecting electoral results.[1] The question then arises as to whether exposure to fake news actually affects voting behaviour or if it only reinforces predetermined political beliefs. In this paper, we attempt to provide an answer to this question by collecting regional data on voting preferences and fake news exposition from the Italian autonomous provinces of Trentino and Alto Adige/Südtirol, and analysing it in a quasi-experimental setting.

We intend to follow the path laid out by Allcott and Gentzkow (2017), and study fake news as intentionally fabricated information characterised by its politically charged content. Fake news in Italy has enjoyed quite a large exposure in anticipation of the 2018 elections and, in a recent report from Il Sole 24 Ore, 5 Star Movement (M5S) and Lega Nord (LN or Lega) voters have been estimated believing (52% and 49%, respectively) and sharing (22% and 11%) fake news in much higher proportions than individuals who did not vote for a 'populist' party.[2] As empirical evidence points at increasing internet and social network usage in Italy growing across all ages and educational cohorts, with middle aged and lower educated individuals experiencing the greatest increase (Istat-Fub, 2018), it is fair to argue that many more Italians had access to false information in 2018 than in 2013, when the previous general elections took place.

Whether the diffusion of fake news has affected electoral outcomes is then a legitimate question that deserves to be answered. In this regard, Italy – and the Trentino-Alto Adige/Südtirol region, specifically – presents the ideal research setting for the study of the phenomenon. First of all, the multiparty nature of the Italian political system allows us to analyse electoral outcomes with much more granularity than in countries with majoritarian systems. Moreover, the last two elections in 2013 and 2018 have been characterised by strong electoral performance for political parties whose platform has been often described as 'populist': the M5S, which topped the poll in each of the two general elections it contested, and the more well-established Lega, whose gradual transformation from a Northern autonomist party to a country-wide 'national-populist' outfit was met by unprecedented success in 2018. Also, as discussed in Campante et al. (2017), part of the growth of these forces can be attributed to a surge in

[1] See Parkinson, H. J. (2016); Click and elect: how fake news helped Donald Trump win a real election; The Guardian; available at: https://www.theguardian.com/commentisfree/2016/nov/14/fake-news-donald-trump-election-alt-right-social-media-tech-companies; last accessed: 15 February 2019

[2] Il Sole 24 Ore (2018); Fake news: quando le bugie hanno le gambe lunghe; ilsole24ore.com; Available at: http://www.infodata.ilsole24ore.com/2018/05/04/fake-news-le-bugie-le-gambe-lunghe/; last accessed: 15 February 2019

participation of previously excluded voters, fostered in turn by increased access to broadband internet connections.

Several areas of Italy are home to significant linguistic minorities. In Trentino-Alto Adige/ Südtirol, in particular, both German and Italian-speaking communities are represented in sizeable numbers. Remarkably, studies show that only a small portion of the local population is functionally bilingual (Ebner, 2016), which suggests a certain degree of separation between linguistic communities. We intend to exploit the language differences across the two *province autonome* that make up the region (Alto Adige/Südtirol, in English 'South Tyrol', is majority-German, and overwhelmingly so in rural areas, while Trentino is in practice monolingually Italian) as an exogenous source of variation in exposure to fake news.

Controlling for the electoral trends between the Italian and German-speaking population, and considering how fake news has been known to spread through channels filtered by an individual's 'echo chambers' (Allcott and Gentzkow, 2017, Boutyline and Willer, 2016), we then advance the hypothesis that German-speaking citizens have been exposed to fake news concerning the Italian elections in a lower magnitude when compared to their Italian-speaking counterparts.

We assume – and this assumption will be tested empirically – that the German-speaking population in Trentino-Alto Adige/Südtirol, while comparable to its Italian-speaking counterpart in terms of economic and demographic conditions, is exposed to a peculiar filter bubble where exposure to fake news concerning Italian politics is limited. Indeed, in line with the approach of Allcott and Gentzkow (2017), fake news sources may be assumed to hold economic or agenda-driven incentives to spread false information. We believe that these incentives are not met in this case: from such a small population as the German speaking population in South Tyrol, website accesses may not generate enough advertising revenue, and the impact on national elections may also be considered negligible as well. In this way, after controlling for electoral trends specific to each language group, we believe we are able to assess the impact of fabricated news over electoral outcomes. This is, to the best of our knowledge, the first attempt to study this phenomenon using a quasi-experimental methodology.

The remainder of the paper is organised as follows. Section 2 provides an overview of the literature on misinformation, filter bubbles and electoral outcomes. Section 3 covers background information on the demographical features and political traditions of the Trentino-Alto Adige/Südtirol region, along with contextualising the spread of fake news within the 2018 Italian general elections. Section 4 describes our data sources, and Section 5 develops a text mining methodology for measuring populism. Our econometric model and results are contained in Sections 6 and 7, respectively. Section 8 concludes and proposes a simple theoretical model assisting us in understanding and discussing our results.

2 Literature review

Due to the nature of the phenomenon and the growing interest from both academics and policy-makers, the field of empirical research on online misinformation has witnessed a considerable growth. Yet a number of challenges linked with developing a successful research design leave many aspects of this phenomenon unexplored.

Fake news and the ability to correctly recognise their mendacity in correlation with prior political beliefs have been studied in a seminal paper by Allcott and Gentzkow (2017). Here the authors also investigate the effects of exposure to fake news – assessing recall rates through employment of placebo headlines – and the nature of the sources of false information, developing a database of fake news and a post-election survey in the process.

Other empirical works include Guess et al. (2018), who uncover evidence of heterogeneous effects conditional on partisan beliefs,[3] estimating that one out of four Americans happened to have been visiting fake news websites in the weeks leading to their interview, and finding out that fact-checking websites have a limited reach on fake news consumers. Indeed, long before the rise of a 'fake news epidemics', Nyhan and Reifler (2010) found evidence that corrections are rarely successful in rectifying misconceptions among targeted groups. In this context, the work of Roozenbeek and van der Linden (2018) – where educational instruments that would allow the public to discern between true and false information are tested in a controlled experimental setting – is certainly worth mentioning.

In any case, any attempt to study the effects of online misinformation cannot disregard the work of Sunstein (2002, 2018) and Pariser (2011) on online echo chambers and filter bubbles, which ultimately highlights how the spread of misinformation is facilitated by online interactions taking place in extremely personalised social media environments. Past experiments on non-online interactions already showed how peer-effects in homogeneous groups can affect political beliefs and the perception of reality (see Schkade et al., 2007). More recently, Boutyline and Willer (2016) find that people sharing conservative or extreme political views tend to seek reaffirmation in their views associating themselves with similar individuals in their online interactions, while Del Vicario et al. (2016) show how false information is propagated through homogeneous online clusters – echo chambers, essentially – each characterised by different cascading dynamics.

Other efforts, such as Törnberg (2018) and Azzimonti and Fernandes (2018), also attempt to model the spread of fake news and its effects on political polarization, providing a much needed theoretical anchoring to the study of misinformation. Finally, a number of studies has also focused on tracking the patterns of diffusion of misinformation (such as Shin et al., 2018, and Allcott et al., 2019), and on developing algorithms for detecting fake news (Shu et al.,

[3] As most of the traffic on fake news websites appears to originate from individuals with extreme conservative views.

2017). So far, the permeability of the social network bubbles between language groups living in the same area has not yet been investigated, and we intend to further explore this topic in the present work.

While most of the aforementioned research points at an evident connection between fake news and populism, the assessment of a causal relationship between misinformation and voting behaviour has proven rather difficult. Due to a number of endogeneity issues, exacerbated by the difficulties in finding a proper quasi-experimental setting, it is still unclear how much online misinformation contributed to the rise of populism, and how much populist sentiments are linked with the inability to recognise false articles. A first step in this direction has been made in Gunther et al. (2019), who find that, among Obama voters in the 2012 US presidential elections, those who believed in fake news were more likely to vote for Trump in 2016. These results are far from definitive: reverse causality may still have led to an overestimation of these figures since, as reported in Weigel (2016),[4] there is a number of confounding factors which may have led these people to switch voting platform, and which we believe may be correlated with the ability to recognise false information. The necessity to control for these unobservable trends and characteristics, along with the bias produced focusing only on the fraction of the population who voted for Obama, leave the question open for further research.

A number of works relied on quasi-experimental approaches to address empirically political economy questions. Our research design draws inspiration from this literature and in particular from the works of Madestam et al. (2013), and of Martin and Yurukoglu (2017). The first studied the effect of political protests on electoral outcomes in the US, taking advantage of the random variation of rainfall as an instrumental predictor for participation in Tea Party rallies. The second estimated the effect of exposition to slanted news on television over voting behaviour by instrumenting channel positions as a predictor for viewership. In a similar fashion, Durante et al. (2019) studied the effect of entertainment television on electoral outcomes in Italy, exploiting the staggered introduction of Berlusconi's commercial TV network in the country since the 1980s. They found that, while municipalities exposed to entertainment TV displayed higher support for Berlusconi's party, in 2013 such support shifted to the populist Five Stars Movement (M5s), suggesting that exposure to entertainment TV made voters generally more supportive of populist parties. Our econometric model is also certainly inspired by diffence-in-difference methodologies, and in this regard the notorious design from Card and Krueger (1994) has stood as an inspiration.

[4]Weigel, David (2016); How voters who heavily supported Obama switched over to Trump; The Washington Post; Last accessed: 02/11/2018; Available at: https://www.washingtonpost.com/politics/how-voters-who-heavily-supported-obama-switched-over-to-trump/2016/11/10/65019658-a77a-11e6-ba59-a7d93165c6d4_story.html?noredirect=on&utm_term=.51b652a76af2

3 Background

3.1 Trentino and South Tyrol: political and sociolinguistic background

In the context of the 2018 Italian general elections, the *provincia autonoma* of Bolzano/Bozen ('South Tyrol' in the following[5]) serves as a natural experiment. Once part of the Austrian Empire and annexed to the Italian unitary state as late as 1919, this Alpine province, home to a little more than 500,000 inhabitants (527,750, Astat, 2017), is the only part of Italy where a sizeable majority of the population is not Italian-speaking.

In 2011, 69.7% of South Tyrol's population declared German as their first language (Astat, 2011). Italian speakers represent slightly less than one quarter of all inhabitants (118,000 people). They are concentrated in the largest municipalities – including the provincial capital, Bolzano/Bozen, where they make up almost three quarters of the population – and in the southernmost part of the province. Conversely, in smaller municipalities, a virtual totality of the population (usually around 95-97%) declares allegiance to the German linguistic group. A third minority language, Ladin, also has legal status in the province: in particular, it is widely spoken as a first language in a number of small municipalities in the east.

Like South Tyrol, Trentino was integrated into the Italian unitary state only after the First World War: before that, it used to constitute the southern third of the Austrian county of Tirol. Other than shared history, the provinces feature striking similarities. Trentino, with 540,000 inhabitants (Istat, 2018) is only slightly more populated than South Tyrol; both provinces are highly rural, with a large share of the population living outside of the few mid-sized urban centres,[6] scattered across hundreds of very small municipalities. Most crucially, both enjoy a large degree of self-government compared to other Italian local authorities, with the *provincia autonoma* having significant legislative, fiscal and budgetary autonomy.[7]

[5]Due to its complex history and strong nationalistic currents on either side of the language barrier, place name choices have long been a matter of contention in this territory. Owing to common usage in English-language sources, in this study we will adopt the name South Tyrol and the demonym South Tyrolean to refer to the territory and population of the officially tri-lingual *provincia autonoma* di Bolzano – Alto Adige (IT)/Autonome Provinz Bozen – Südtirol (DE)/Provinzia autonoma de Bulsan – Südtirol (Ladin). In Italian-language sources, 'Alto Adige', a name with a strong historical association with Italian nationalism, and the demonym 'altoatesini' are relatively more common. Together with the provincia autonoma of Trento (commonly referred to as Trentino), South Tyrol is part of the Trentino–Alto Adige/Südtirol region. In the following, we will adopt the official, bilingual denomination in full when referring to the region as a whole. For municipalities and other place names in South Tyrol we will generally use the official denomination (bilingual, Italian name first, separated by a slash, e.g. in Bolzano/Bozen). For the sake of simplicity, in our primary data municipalities are listed under their Italian-language name.

[6]Istat (2015); Principali dimensioni geostatistiche e grado di urbanizzazione del Paese; istat.it; Last accessed: 25/05/2019; Available at: https://www.istat.it/it/archivio/137001

[7]For an overview of the competencies of the *province autonome*:
Trentino: Consiglio della provincia autonoma di Trento; "Le competenze legislative secondo lo statuto"; consiglio.provincia.tn.it; Last accessed: 25/05/2019; Available at: https://www.consiglio.provincia.tn.it/istituzione/l-autonomia/il-regime-delle-competenze-legislative-e-amministrative/Pages/il-quadro-delle-competenze-legislative-secondo-lo-statuto.aspx
South Tyrol: Amministrazione Provincia Bolzano; "Competenze e finanziamento dell'autonomia"; pro-

The main divergence between the two provinces is, therefore, language: with limited exceptions, Trentino is essentially monolingually Italian. This is, to some extent, an intended effect of the post-Second World War settlement that confirmed the sovereignty of Italy over the area, fending off separatist tendencies among the German-speaking population. The De Gasperi-Gruber agreement between Italy and Austria (1946), which enshrined a set of safeguards for the German community (such as native language education, public sector employment quotas, reversal of assimilationist practices), and still represents the basis for the current institutional setup, soon led to the transfer of a few German-majority municipalities from the province of Trento to that of Bolzano/Bozen. Indeed, subsequent pieces of legislation increased the degree of autonomy of South Tyrol not just from Rome, but from Trento too: tellingly and uniquely in Italy, legislative functions are given to the provincial authority (hence, *provincia autonoma*) and not to the Trentino–Alto Adige/Südtirol region, which is currently little more than a coordinating body.

All of this means that the border between the two provinces is akin to a linguistic frontier, with no meaningful German community to the south of it. Even though there are recognised and protected minority languages in Trentino – the aforementioned Ladin, plus two archaic Germanic dialects, Cimbro and Mocheno – their level of diffusion and their legal status is not comparable to that enjoyed by German (and Ladin) in South Tyrol. Nonetheless, census data shows that a few municipalities in the north-east of the province have a non-Italian (generally Ladin) majority (Servizio Statistica provincia TN, 2012).

South Tyrol acts as an unusual filter bubble in the Italian context, clearly showing evidence of linguistic segregation. *Dolomiten*, the main German-language newspaper, boasted in 2016 a circulation of 42,103:[8] this is more than four times higher than that of its main, local Italian-language counterpart, *Alto Adige*, and exceeds that registered by the most common mainstream newspapers of the country ('Corriere della Sera' and 'La Repubblica') over ten times. Alongside Rai, the national broadcasting operator – which also maintains German-speaking radio and TV programming – a Province-funded broadcaster relays transmissions from surrounding German areas. Although both Italian and German are compulsory subjects for members of both language groups from age six, and graduate-level jobs in public administration require high proficiency in either language, effective bilingualism is not widespread. Past investigation, cited in Ebner (2016), points to a relatively low second language proficiency (L2) of the South Tyrolean population. In particular, less than 10% of 17-18 year old high school pupils of either language group have reached a proficiency of the other language above the C1 level.

The extent of the language divide is emphasised in the political dynamics of the two provinces. While devolution certainly plays a part in Trentino's politics, its voting trends,

vincia.bz.it; Last accessed: 25/05/2019; Available at: http://www.provincia.bz.it/politica-diritto-relazioni-estere/autonomia/competenze-finanziamento-autonomia.asp

[8] Accertamenti Diffusione Stampa (ADS) (2016), *Dati medi annuali territoriali per testata: diffusione cartacea Italia*, available at: http://www.adsnotizie.it/_grafici.asp? (Last accessed August 2018)

especially in national elections, do not dramatically diverge from neighbouring areas. Like in much of the Italian north-east, Trentino's dominant force used to be Democrazia Cristiana (DC), which scored over 50% of the votes in every general election until the 1980s. With DC's slow demise, the province has experienced an increase in competition between the centre-left and the centre-right blocs. In fact, leaving 2018 aside – a year that has seen a decisive win for the centre-right parties, Lega in particular – the centre-left bloc fared relatively well in Trentino compared to the rest of Triveneto. This may be attributed to an increased prominence of localist Catholic-centrist parties, a force to be reckoned with in provincial elections since the 1990s, which were generally aligned with centre-left parties at the national level (either running in coalition or backing nationwide lists). Indeed, no localist pro-autonomy party has ever reached double digits in general elections in Trentino.

In South Tyrol, conversely, the situation is that of a parallel party system segmented across linguistic divides, to an extent that is not experienced anywhere else in Italy. Since the birth of the Italian Republic (1946), South Tyrolean politics has been marked by the dominance of the ethnic German South Tyrolean People's Party (Südtiroler Volkspartei, 'SVP' in the following). A typical catch-all party, Catholic-based, featuring both conservative and social-democratic wings, SVP obtained a majority of the votes cast in the Province in every election between 1948 and 2008, scoring landslides (often more than 90% of the vote) in several of the purely German-speaking municipalities. A remarkable exception is the area of Bolzano/Bozen, where the SVP has rarely exceeded 25% of the votes. Nevertheless, the SVP can hardly be considered a populist platform and, in its history, has always linked itself with traditional political forces such as the *DC* before 1992 and the *Partito Democratico* (PD) and its predecessors afterwards.

In general elections, SVP has mostly faced provincial-wide competition from Italian parties only. In the last decades, however, several regionalist and separatist parties, for long a relevant force within local and provincial elections, have also ran. In 1996 – an election in which, due to a quirk of the electoral system then in force, the SVP logo did not appear on the ballot on its own right – a conservative splinter got almost 20% of the votes. In 2008 and 2013 – elections that were fought with a proportional-based system, *Die Freiheitlichen* (DF), a radical-right separatist party akin to Austria's Freedom Party, received 9.5% and 15% of the vote, respectively. Consequently, in parallel with the electoral decline of the centre-left in Trentino, the provincial vote of the SVP has been progressively eroded, now stopping well short of 50% overall.

In 2018, however, DF did not field candidates in the national election, admittedly because the electoral law introduced for that year left very slim chances to win seats for German competitors to the SVP.[9] Left as the only German party, the SVP scored more than 60% of the votes in majority-German constituencies, but only 13.4% in Bolzano/Bozen proper. Turnout

[9] Alto Adige (2017), *I Freiheitlichen*, altoadige.it, available at: http://www.altoadige.it/cronaca/bolzano/i-freiheitlichen-1.1478219

in German-speaking areas also registered a remarkable drop, from over 80% to slightly more than 60%.[10] It can thus easily be assumed that a rather large share of potential SVP voters opted to abstain in this occasion.

Figure 1 displays the share of Italian speaking voters for each municipality in the Trentino-Alto Adige/Südtirol region,[11] along with the variation in populist vote between the two elections.[12] While populist vote has been increasing across both linguistic groups, it is evident that populist platforms have had a greater appeal towards Italian-speaking voters. This trend seems to persist even in the areas where the two groups are more evenly mixed, and the increase appears particularly marked in Italian-speaking rural municipalities. Of course, there are confounding factors to control for, and without information on exposition to misinformation the connection between fake news and populism growth cannot yet be quantified. Still, these figures prove that linguistic filter bubbles have had differential effects with regard to the rise of populist platforms.

[10] Italian Ministry of Interior, Archivio storico delle elezioni, Dipartimento per gli Affari Interni e Territoriali. Available at: https://elezionistorico.interno.gov.it/

[11] Source: authors' calculations from 2001 and 2011 census data; available at https://astat.provincia.bz.it/downloads/mit38_2012.pdf; and http://www.statistica.provincia.tn.it/statistiche/societa/popolazione/; last accessed: March 9, 2019

[12] Source: Italian Ministry of Internal Affairs, Historical Archive of Italian Elections; available at: https://elezionistorico.interno.gov.it/index.php?tpel=C; last accessed: March 9, 2019

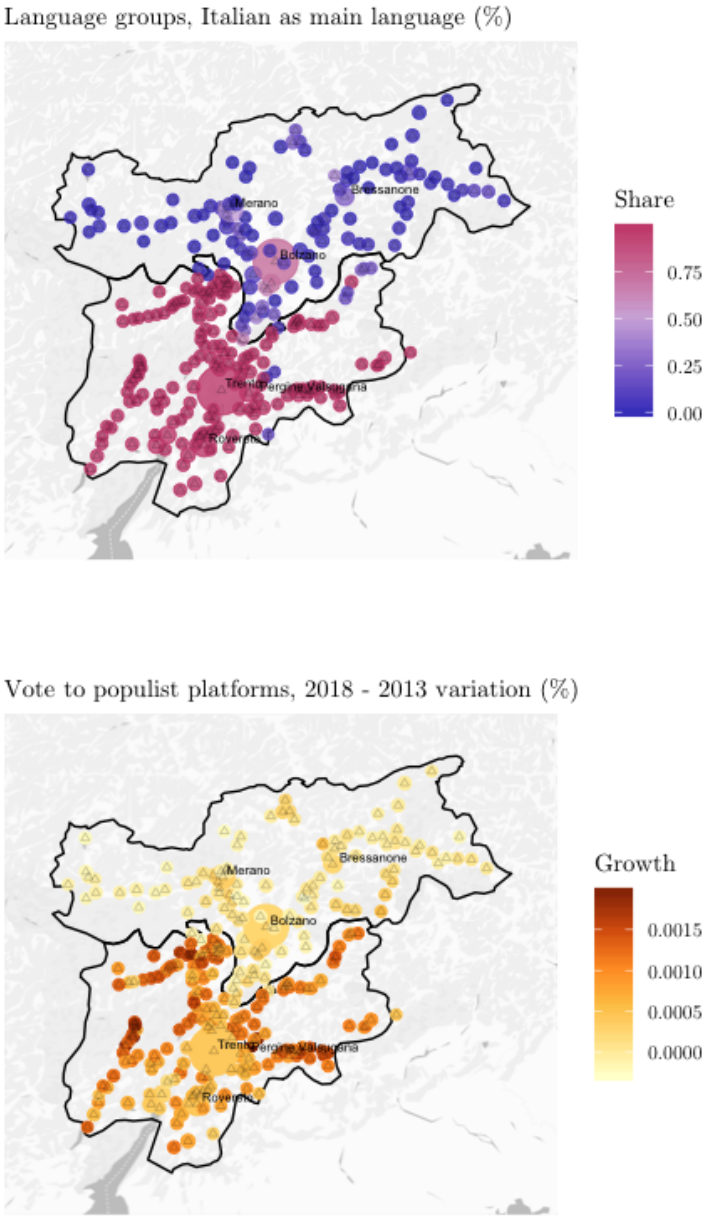

Figure 1: Language groups shares and growth of populist parties (2013-2018) in Trentino-Alto Adige/Südtirol. The size of the circle around each municipality indicates average electorate size between the two periods. Language groups shares are computed for the period in-between the two elections, estimated through a non-parametric linear interpolation of Aitchison's (1986) log-ratios for the 2001 and 2011 census data shares. Populist scores are computed by multiplying votes for each party for the relative score obtained through the 'assertive' dictionary method, summing all resulting scores by municipality, and then dividing for the number of voters.

3.2 Fake news in the 2018 Italian general elections

According to several journalistic[13] and institutional[14] sources, the campaign period leading to the 2018 Italian general election saw a remarkable spread of 'fake news', i.e. entirely baseless news stories published by non-institutional outlets, on social media.

These claims have been easily substantiated by a recent investigation performed by an Italian news channel[15]. Tracking a limited set of politically-charged keywords via a content analysis tool (footnote: Buzzsumo), it has been found that among the top 100 articles in Italian for social media engagement, 5 were hoaxes, while another 10 were judged even 'highly problematic' by the authors. While the 'purely false' news has a predominantly 'anti-establishment' character (and, by extension, anti-PD, the incumbent party), many of the 'problematic' (i.e. misleading) news in the sample focuses on immigration. According to an extensive report (Giglietto et al., 2018), immigration and security were dominant themes on social media during the run-up to the 2018 Italian General election.[16]

While few in number, the exposure reached by these hoaxes was highly significant. The second-most shared news in the SkyTg24 database, published on the day before the election and consisting of an entirely unsubstantiated report of voter fraud in Sicily to the advantage of PD, received more than 140,000 interactions, mostly on Facebook. As pointed out by some observers, the news seemed to target a public which was friendly to the 5 Star Movement, and was widely shared among M5S supporters (including an MP).

In addition, the name of the originating website, the now offline 'ilfatto.it', is a potentially misleading reference to *Il Fatto Quotidiano*, a daily newspaper very popular among M5S voters. This 'spoofing' tactic of mirroring more reputable news sources was exploited by a number of partisan outlets, that occasionally managed to create an engagement on social media comparable to their 'genuine' peers. Unsurprisingly, their URL also appeared in 'blacklists' compiled by the most popular fact-checking websites in the country (Butac.it, Bufale.net).[17]

Giglietto et al. (2018) provides a detailed classification of news sources based on partisanship of their news content, evidencing that a vast majority of 'non-institutional' web-

[13] See Buzzfeed News (2017); One Of The Biggest Alternative Media Networks In Italy Is Spreading Anti-Immigrant News And Misinformation On Facebook; available at: https://www.buzzfeed.com/albertonardelli/one-of-the-biggest-alternative-media-networks-in-italy-is; last accessed: 25 May 2019

[14] See Autorità per le Garanzie nelle Comunicazioni (2018)

[15] Bruno, Nicola (2018); Satira e fake news: gli articoli più condivisi delle Elezioni 2018; Available at: https://tg24.sky.it/politica/2018/03/08/fake-news-elezioni-2018.html

[16] This was also influenced by a real, contentious news story involving the alleged murder of an Italian teenager in Macerata, Central Italy, which was followed by a racially motivated attack against a group of Africans in the town. See also: The New York Times (2018); This Italian Town Once Welcomed Migrants. Now, It's a Symbol for Right-Wing Politics; available at: https://www.nytimes.com/2018/07/07/world/europe/italy-macerata-migrants.html; last accessed: 25 May 2019

[17] These black lists are available at the following links: https://www.butac.it/the-black-list/ and https://www.bufale.net/the-black-list-la-lista-nera-del-web/. Last accessed: 29 May 2019.

sites feature some form of bias towards of the largest Italian populist parties, Lega and M5S. Crucially, comparable biased sources supporting pro-establishment (and smaller anti-establishment) parties captured much less social media attention than the pro-M5S and Lega networks.

The study however stops short of establishing a link between the spread of false information in the electorate and the support for Italian populist parties and their policy stances. Nonetheless, a recent investigation by Avaaz[18] provided evidence in support of this link, at least as far as Facebook is concerned. The report uncovered an extensive network of Facebook pages and fake accounts (with a reach of millions of interactions) that, in blatant violation of the platform's terms of use, pushed misinformation in support of Lega, M5S and other anti-establishment and fringe causes – including antisemitism and racism.[19]

As a complement to these findings, we provide further evidence for this anti-establishment bias by analysing the topics occurring in a sample of fake news propagated between March 2016 and March 2018, the month of the elections. We do so by analysing entry metadata scraped from an independent Italian debunking website (Butac.it) and retrieving from it the text of all fake news headlines it reported.[20]

Using a simple text mining technique on fake news headlines,[21] we search for recurring mentions of parties and leaders of incumbent/establishment and challenging/anti-establishment platforms, along with terms denoting topics which have been predominant during the 2018 electoral campaign: namely, immigration, foreign policy and the European Union, and vaccinations.[22]

Figure 2 plots the relative frequency of these topics over a two-month period, showing that the challenging platform received much less attention than any other topic. Moreover, fake news' mentions to anti-establishment parties and leaders were overall less frequent than mentions to incumbent ones, providing additional descriptive evidence in support of fake news' anti-establishment bias.

[18] La Repubblica (2019); Facebook chiude 23 pagine italiane con 2.4 milioni di follower: diffondevano fake news e parole d'odio; available at: https://www.repubblica.it/tecnologia/social-network/2019/05/12/news/facebook_chiude_23_pagine_italiane_con_2_4_milioni_di_follwers_diffondevano_fake_news_e_parole_d_odio-226098817/; last accessed: 13 May 2019

[19] As a result, in May 2019 Facebook took down 23 of the pages reported by Avaaz. While some of these pages explicitly portrayed themselves as unofficial supporters of the two parties, others used a more subtle approach. The most popular page in this subset, 'I Valori Della Vita' (1.5m followers), ostensibly a lifestyle page, was actually part of a bigger network, sharing in a coordinated manner content from a right-wing, pro-Lega news site.

[20] The website was scraped in January 2019.

[21] We use a dictionary technique which computes a score for each topic based on the occurrence of key terms related to each topic in the textual database. For this purpose, dictionary techniques require the creation of list of key words that relate to a specific topic and that the algorithm will search in the text subject to the analysis, that is the fake news headlines in our case.

[22] The text bags used to construct these indicators are based on the available fake news headlines. The codes and data used in Figure 2 (including all text bags) are available in the online data archive.

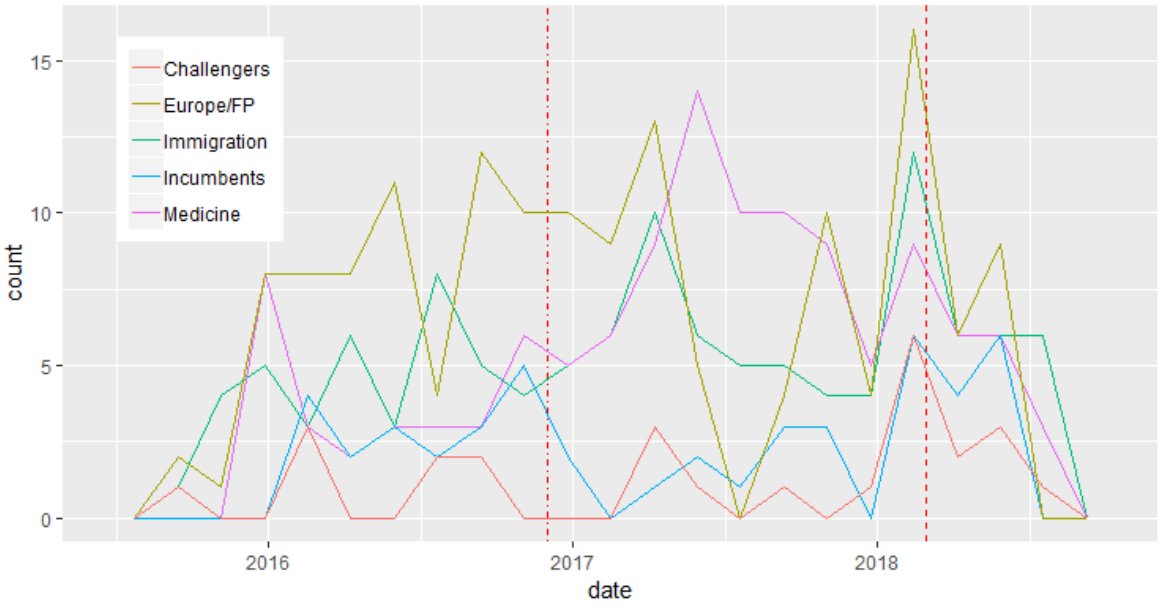

Figure 2: bimonthly frequency of pieces of fake news by topic. The dot-dashed and dashed red lines indicate the date of the 2016 Constitutional Referendum and the 2018 General Elections respectively.

In addition, it should be noted that, when a false fact pertains to the challenger, the topic does not necessarily indicate the slant of the piece of misinformation: while most headlines pertaining to the incumbents, the EU, immigration and vaccinations show an obvious anti-establishment bias, the same cannot be said for the challenging platform. On the contrary, most of the few pieces of misinformation concerning the Five Star Movement and the Lega Nord take a favourable and supportive stance with regard to these parties, meaning that the actually damaging headlines are even sparser than what the figure would suggest. Indeed, within the time window we analysed, only 6 among all headlines reported by Butac.it could be said to have an anti-M5S or anti-Lega slant, while 2 of them were directed against the neo-fascist parties Casapound and Forza Nuova.[23] It is also interesting to note that each topic experienced differential increments in their frequency in proximity to the 2016 Constitutional Referendum and the 2018 General Elections, suggesting that the propagation of fake news may be driven by electoral purposes.

While it is not our intention to assess whether the spread of 'fake news' was an integral part of Italian populist parties' strategy and campaign, in this paper we will assume, on the grounds of overwhelming empirical evidence, that misinformation and outright hoaxes have been exploited by partisans of Lega and M5S on social media in the run-up to the 2018 general

[23]The 4 headlines with an anti-M5S stance are available at these links: https://www.butac.it/di-maio-emiliano-e-matera/; https://www.butac.it/sindaco-grillino-attacca-zanardi/; https://www.butac.it/sahaja-yoga-e-la-candidata-del-m5s/; https://www.butac.it/piccole-perle-di-facebbok-i-soldi-del-m5s/. The other 2 headlines concerning the Lega are here available: https://www.butac.it/salvini-i-giovani-padani-e-quelli-del-mezzogiorno/; https://www.butac.it/democratica-salvini-calabresi/. Last accessed: 29 May 2019

election, and that this effort was much more successful than that of their pro-establishment counterparts *in terms of social media exposure*. The goal of our study is to assess precisely whether this social media success causally translated to shifts in voting behaviour.

4 Data

Due to the lack of individual-level micro-data on fake news exposure and voting preferences, we used municipalities, where much more information is available, as the principal unit of analysis of this study. Our research needs required us to collect information on voting preferences, exposure to fake news, and important socio-economic features such as linguistic group shares, income levels and internet coverage. Moreover, information on populist stances for each party running for election was also necessary. To address these needs, we employed a number of different sources in the construction of our final dataset.

As in-between the two elections some municipalities were suppressed and new ones were formed from their merger, a few adjustments had to be made. For all suppressed municipalities from 2013, values for the corresponding 2018 entities were imputed either by summing the totals or, in the case of per capita figures, by averaging population-weighted sum of the now-defunct municipalities.

The present section provides an overview on how these sources were gathered and harmonised. Further information on how misinformation data was collected and processed is available in Appendix A. All data sources and codes used to construct our final dataset and perform our analyses are available in the online data archive.

4.1 Electoral data

Official municipality-level data on general election results from the Trento and Bolzano/Bozen provinces are obtained from the Italian Ministry of Internal Affairs,[24] where the complete election history of the Italian Republic is available. We extracted information on electoral results from the 2013 and 2018 general elections.

Other than indicating how many votes each party received in every election, the data set also includes a number of auxiliary variables disaggregated by municipality. Among these, electorate size and total number of abstentions or invalid votes were of particular interest for our research. While voting outcomes have been later used to construct our dependent variable, these other variables have found important applications in our work. Electorate size, other than being used for weighting other variables, has proven an important predictor of populist preference. The same can be said for non-voting behaviour or abstention, which has also been

[24]Italian Ministry of Internal Affairs, Historical Archive of Italian Elections; available at: https://elezionistorico.interno.gov.it/index.php?tpel=C; last accessed: March 9, 2019

used to control for potential voters of the Die Freiheitlichen party across the German-speaking subgroup.

4.2 Socio-demographic and internet connectivity data

The proportion of language groups by municipality is a key variable in our identification strategy. Language group shares for the region can be obtained through census data.[25] However, as census data for Italy is only released once every 10 years, the latest figures available date back to 2011. To compensate for differing trends in population growth across language group which may have affected our analysis, we perform a small and simple adjustment by interpolating the figures from 2001 and 2011 to predict the group-specific shares for the years 2013 and 2018, using Aitchison's (1986) log-ratio transformation to preserve the compositional form of the data. While the great majority of locations in Trentino are overwhelmingly Italian-speaking, language groups shares can vary significantly within South Tyrol: consequently, this indicator allows us to address the correlations between language groups and exposition to misinformation with much more precision than a binary discriminant between the two provinces would allow.

To reconstruct per capita income by municipality, we used tax data from the Italian Ministry of Economy and Finances for the years 2012 and 2017.[26] Those calculations are based on self-declared taxable income from the *Imposta sul reddito delle persone fisiche* (IRPEF), and per capita figures are already available by municipality. The connection between income and voting preference is not unambiguous, but has been discussed extensively in Galbraith and Hale (2008), Lewis-Beck and Nadeau (2011), and Hersh and Nall (2015), among others. Income differences can also arise across language groups, motivating our choice to control for them in our econometric specification.

Two different sources were used, instead, to construct statistics for broadband internet connectivity in 2013 and 2018. As mentioned earlier, Campante et al. (2017) found that increases in internet connectivity have had an effect on voting preferences in Italy, suggesting that a measure for connection latency should be used as a control in our econometric model or even as an alternative to our language group instrument. Also, as Italian regions – including the two *province autonome* of Trentino and Alto Adige/Südtirol – possess considerable autonomy in the implementation of broadband infrastructures, this indicator then plays a vital role in controlling for differences in connectivity arising from staggered local legislation.

'Digital divide' in Italy, defined as the share of households not covered by broadband connection, was originally covered by a set of municipality-specific indicators released by the

[25] Province-specific census data on language shares is available at: https://astat.provincia.bz.it/downloads/mit38_2012.pdf; and http://www.statistica.provincia.tn.it/statistiche/societa/popolazione/; last accessed: March 9, 2019

[26] Available at: https://www1.finanze.gov.it/finanze3/analisi_stat/index.php?tree=2013 and https://www1.finanze.gov.it/finanze3/analisi_stat/index.php?tree=2018; last accessed: March 31, 2019

'Agenzia per la Coesione Territoriale'.[27] This indicator, which defined any landline connection whose speed exceeded 2Mbps as broadband, was then subtracted from unity to compute the share of low latency connections by municipality. As these publications were discontinued, the collection duty was then moved to the 'Autorità per le Garanzie nelle Comunicazioni' (AgCom), which has released internet connectivity indicators for 2018.[28] The new variables released by AgCom were slightly different from the previous ones, but we managed to reconstruct a fully comparable digital divide indicator using the available information on the number of households with a landline connection faster than 2Mbps.

4.3 Social media data

Information on fake news exposition and on the social media campaigns of Italian parties was all scraped from Facebook. Our fake news exposition indicator shows, for each municipality and for each year of election, the estimated number of Facebook likes being held by all Facebook fan pages that are known to spread politically-charged misinformation.[29] Due to the unavailability of granular data on the spread of each piece of false information, we decided to focus our attention on their 'disseminators', measuring – in a given municipality – the effect of each unitary increase in their social media following on aggregate electoral preferences. Our estimates will then include both the 'intensive' and 'extensive' margin of fake news exposition, under the implicit assumption that individuals who 'liked' these pages also play an active part in spreading misinformation. The lengthy data collection and estimation process behind the construction of this variable is reported in Appendix A.

Our research question also required us to transpose categorical voting decisions into a continuous scale measuring affinity to populist discourse by municipality. The construction of such an indicator required access to propaganda content used over the course of the two electoral campaigns preceding the 2013 and 2018 elections. We turned, again, to Facebook, from which we scraped all posts from running parties and their leaders, covering all three months preceding the elections in 2013 and 2018, and coinciding with the beginning and the end of each electoral campaign. A more detailed description of the data and calculations leading to the production of our final indicators is provided in the next section.

[27] Available at: http://old2018.agenziacoesione.gov.it/it/arint/OpenAreeInterne/index.html; last accessed: March 31, 2019

[28] Available at: https://maps.agcom.it/#; last accessed: March 31, 2019

[29] Likes from all pages which appeared in black lists compiled by debunking websites Butac.it and Bufale.net (as discussed in section 3.2) were estimated. See Appendix A for the full list of pages used in the process.

Table 1: Summary statistics by province and year

	(1) Bolzano/Bozen			(2) Trento		
	2013	2018	Total	2013	2018	Total
Populist score (total)	5.972	6.526	6.249	4.123	5.899	5.011
	(13.376)	(15.823)	(14.621)	(12.074)	(15.465)	(13.883)
Exposed to fake news	148.424	376.876	262.650	199.858	521.874	360.866
	(869.421)	(2366.432)	(1782.496)	(1206.976)	(3459.335)	(2592.059)
Broadband connections	2469.032	3007.955	2738.494	1987.041	2217.626	2102.334
	(7542.113)	(7661.528)	(7590.388)	(6763.404)	(7030.532)	(6889.395)
Italian speaking voters	879.493	907.630	893.562	2173.312	2226.701	2200.007
	(5321.100)	(5435.457)	(5366.947)	(6837.192)	(6963.666)	(6890.933)
Income per capita	18721.313	20692.652	19706.982	17193.012	18434.036	17813.524
	(2534.504)	(2728.883)	(2807.309)	(2188.215)	(1891.333)	(2134.696)
Electorate size	3245.284	3332.853	3289.069	2276.602	2330.716	2303.659
	(7548.803)	(7683.594)	(7600.120)	(6986.644)	(7112.244)	(7039.726)
Abstentions and invalid votes	677.569	1345.802	1011.685	524.455	649.989	587.222
	(1668.889)	(2502.587)	(2148.627)	(1499.638)	(1966.976)	(1747.624)
Observations	232			352		

*N*otes: Mean coefficients, standard errors in parentheses.

5 A text mining approach to measuring the populist content of parties

As described in Section 3, descriptive evidence suggested that fake news benefits populist parties. In order to investigate the electoral impact of fake news, the dependent variable of our empirical model will hence be represented by the electoral performance of these platforms. Nevertheless, defining the degree of populism of each party in an objective manner is not a trivial exercise.

To overcome the issue of objectivity, we apply text analysis to the Facebook posts of political parties and leaders that ran in the 2013 and 2018 Italian elections. In doing so, we define populists as those parties that (1) share an anti-elite rhetoric (Albertazzi and McDonnell, 2008; Pauwels, 2011; Rooduijn and Pauwels, 2011; Kaltwasser et al., 2017) and that (2) tend to adopt a particularly emotional language in their campaign (Taggart, 2000; Rooduijn, 2014; Caiani and Graziani, 2016; Bischof and Senninger, 2018). These two elements are generally complementary and do not exclude one another. As found by Guiso et al. (2017), in fact, the supply of populist parties tends to be higher where disappointment with traditional parties is greater, suggesting a correlation between the rise of populist parties and an emotional narrative against incumbent politicians perceived as part of the establishment.

Following these considerations, we introduce a new methodology to measure the degree of populist rhetoric in each party based on text analysis. In particular, we create two dictionaries that capture (1) an anti-establishment rhetoric and (2) a more emotional or assertive tone, and match them with the text produced on Facebook by all the Italian parties that ran in the 2013 and 2018 elections and their leaders.[30] The first dictionary contains 23 words which capture a populist tone in the Italian political language; for example, the words 'establishment' or 'caste'. The second dictionary is simply capturing the number of exclamation marks as a proxy of an emotional tone.[31] While previous works applied dictionary techniques to measure populism on party manifestos (Rooduijn and Pauwels, 2011; Pauwels, 2011), this paper is the first to our knowledge that use them to measure the degree of populism in the political communication on social media.

We then web-scrape the text of Facebook posts of Italian parties (12.159 posts) and their political leaders (8.164 posts) that preceded the Italian elections in 2013 and 2018. After pre-processing the text,[32] we first match the words of these dictionaries against all Facebook posts collected. Secondly, we sum the number of each match at party level and weight them by the total number of words and the social media interactions per post.[33] Given these transformations, our scores should not be assessed by looking at their absolute values, which are necessarily low, but rather on the relative distance in the scores between parties. Results are displayed in Fig. 3 for both the dictionary (panel on the left) and the exclamation marks (right).

The two indexes display similar trends, further confirming the belief that anti-establishment discourses and emotional tones are generally correlated in political communication. In particular, the Lega clearly features as the most populist party in the election of 2018 in both cases, as a result of a steep increase in populist language from 2013. On the other hand, the 5 Star Movement (M5s) shows a lower degree of populist language in 2018 than in 2013, for both cases. In line with our expectations, the incumbent Democratic Party (PD) displays instead lower levels of populism. The ability of our methodologies to proxy for a populist language is

[30] We decided to focus on the rhetorical aspect of populism rather than on party-specific policy stances. This choice was led by the fact that, while on the one hand defining populism as linked to a specific policy stance is controversial, as it has often been applied to very different contexts (Caiani and Graziani, 2016), on the other hand, populism can be more easily identified by its forms of communication rather than by specific ideological stances (de Vreese et al., 2018).

[31] Previous works in computer science, such as Kumar and Sebastian (2012), found exclamation marks to be a good proxy for the prediction of emotional statements on social media.

[32] In order to allow for matching, we tokenise the text, remove Italian stopwords, punctuation (excluding question marks for the second dictionary), numbers and white spaces and transform all terms in lower case. We also record the number of social media interactions for each post, a duplicate each text-bag in German so that scores can be computed for the SVP and DF parties. With the exception of the text bag capturing assertive tones, which relies on the computation of exclamation marks, we also remove punctuation from the posts.

[33] This practice aims to avoid that long posts inflate the populist score. Intuitively, posts containing larger number of words are more likely to contain a higher number of words that are matched.

further reinforced by the proximity of fringe parties, which are generally anti-establishment, to the Lega and the 5 Stars, such as the far-left Potere al Popolo (PaP) and the neofascist Casapound (CP).

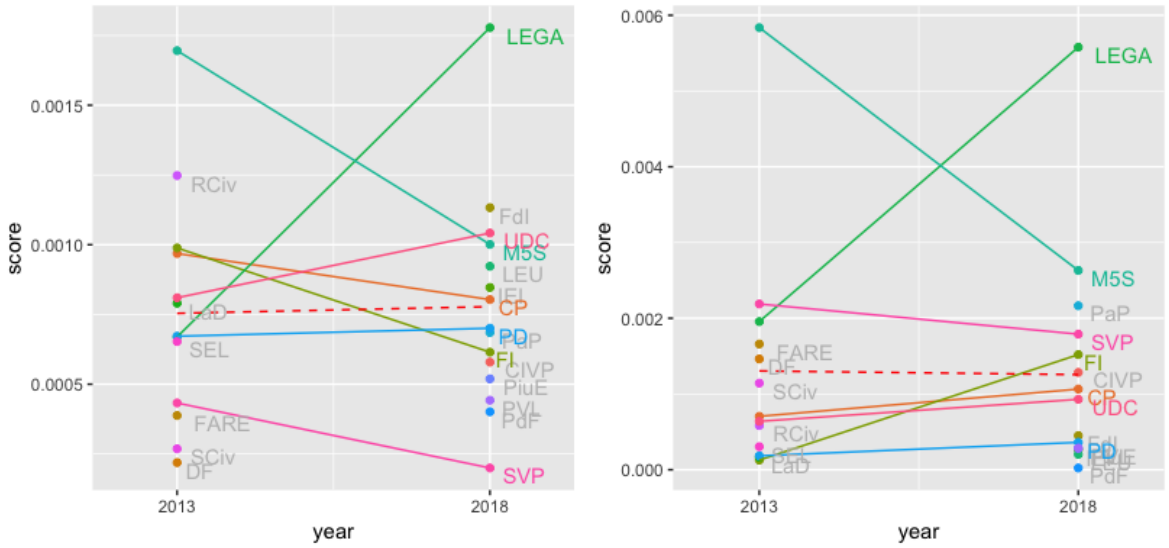

Figure 3: Text analysis scores of social media posts from parties and their leaders during the 2013 and 2018 elections campaigns. The left figure refers to the scores obtained with the populist text-bag, while the right figure computes the same score as the frequency of exclamation marks in the text. Parties in grey have only took part in one of the two elections. The red dashed line refers to the election specific average.

However, such proximity is not constant as, for example, in the first indicator PaP has a similar score to the PD, suggesting that the potential of the indicators might be weaker for minor formations. On the other hand, some of these counter-intuitive results might find an explanation through a deeper study of political dynamics. The clearest example is UDC, a centrist moderate party, which presented particularly high scores in 2018. This shift might be explained by a transformation in the communication style of UDC following the change of coalition in the 2018 election. While in 2013 it ran in coalition with Scelta Civica, a party founded by the former technocratic prime minister Mario Monti, which in fact scores relatively low in both indicators, in 2018 UDC joined a centre-right wing coalition formed by the Lega, Forza Italia (FI) and Brothers of Italy (FdI), which all report high scores in 2018. By inspecting the words matched under the first methodology, terms like 'scandal' or 'shame' appear much more frequently in 2018, reflecting this change in communication style during the two elections.

Despite their relative limitations, these results confirm the widespread categorisation of parties such as the Lega and the M5S as populist. As a further test, we compared our results to the values assigned by political experts in the Chapel Hill Expert Survey (Polk et al., 2017)[34] to a number of parties for the variable 'People vs the Elites', which measures the 'salience of

[34]Data are available at the following link: https://www.chesdata.eu/.

anti-establishment and anti-elite rhetoric'. Text analysis scores display a positive correlation with the anti-establishment indicator based on experts' perceptions, in particular displaying the Lega and the 5 Star Movement among the most populist parties (see Figures 4 and 5 in Appendix C).

6 Econometric model

In an ideal experimental setting, voters would be randomly assigned into two groups: in the 'treatment' group, individuals would be exposed to fake news, while voters in the 'control' group would have access to reputable sources of information only. In this case, differences in voting behaviour would be explained by the assignment to the experimental unit and the average treatment effect would be obtained by the simple difference in means across groups. Such an experiment, however, would be difficult to replicate, as we believe that prior exposition to politically charged misinformation can have lasting effects on voting behaviour that could affect the outcome of the experiment and bias estimates downward. Indeed, random assignment cannot control whether individuals in the control group have already been exposed to fake news in the past. Preventing access to the control group based on exposition to fake news would also invalidate the research, as such exposition could be linked to factors that also influence voting preferences. For this experiment to be unadulterated by environmental factors influencing our object of study, every element of it – treatment and outcome included – has to be detached from their connections with real-world politics, something very difficult to achieve.

This is then one of the few instances when a quasi-experimental setting would be preferable to an experimental one. The South Tyrol province seems like the perfect candidate for a case-study, where individuals are randomly assigned at birth in one of two different linguistic groups. It is not unreasonable to argue that each of these groups is exposed to misinformation concerning Italian national politics in a different way: as mentioned earlier, and following from Allcott and Gentzkow (2017), fake news disseminators may have little to no incentive to produce misinformation across the German-speaking population, whose spread would also be isolated by the linguistic echo chamber each voter belongs to.

There are, however, a couple of complications that have to be taken care of, the first of which consists of the different electoral patterns across the two linguistic groups. As the SVP has consistently proven to be a popular voting choice across the German-speaking population, our estimates would suffer from extreme upward bias when not controlling for previous elections. To account for this issue, we would first employ a difference in difference framework, taking advantage of the different filter bubbles generated across the two linguistic groups in a mixed language area, constructing the treatment group from the Italian-speaking individuals in the sample, and the control group from their German-speaking counterparts. We would

then investigate differences in voting outcomes, exploiting the changes between the 2013 and 2018 elections. In this way, diverging electoral trends can be accounted for. Our diff-in-diff specification would take the following form:

(6.1) $Y_{it} = \alpha + (Z_i \times P_t)\lambda + Z_i\delta + P_t\zeta + X_i^{'}\gamma + e_{it}$

where i represents each municipality and $t \in [2013; 2018]$ stands for each election ranked by their date. The variable Z proxies for the exposition to fake news for voters in municipality i. In this specification Z is represented by the share of Italian-speaking population in one municipality, i.e. by the share of population potentially exposed to fake news due to their linguistic characteristics. P is a dummy that captures time fixed effects for the two elections, whereas X is a vector of covariates controlling for demographic and economic characteristics. As long as assignment to the Italian linguistic group – identified by Z – correctly proxies for exposure to misinformation, and that exposure to fake news took place only prior to the latest elections, then the coefficient λ of the interaction term $Z_i \times P_t$ will indicate the effect of fake news on populist preferences Y in each municipality.

However, as there are reasons to believe that the two language groups are not entirely isolated from each other, we reckon that a simple diff-in-diff specification may not suffice: random assignment to the language group may rather affect the intensity of exposure to fake news. To account for 'permeability' of linguistic filter bubbles, we then turn to a two stage least squares specification where we make use of language group as an instrumental variable, while controlling for both time variant and invariant determinants. We then propose an alternative specification, for $t \in [2013; 2018]$:

(6.2.1) $F_i = \alpha_1 + (Z_i \times P_t)\lambda_1 + Z_i\delta_1 + P_t\zeta_1 + X_i^{'}\gamma_1 + e_{1it}$

(6.2.2) $(F_i \times P_t) = \alpha_2 + (Z_i \times P_t)\lambda_2 + Z_i\delta_2 + P_t\zeta_2 + X_i^{'}\gamma_2 + e_{2it}$

(6.2.3) $Y_{ti} = \alpha_3 + (F_i \hat{\times} P_t)\lambda_3 + \hat{F}_i\delta_3 + P_t\zeta_3 + X_i^{'}\gamma_3 + e_{3it}$

In the present specification, we apply an alternative definition of exposure to fake news, based on the number of likes to Facebook pages that diffuse fake news in municipality i at time t. This exposure to fake news is represented by F in regressions (6.2.1) and (6.2.2), where its value – and the value of its interaction with the year of election in (6.2.2) – is predicted by the instrument Z (indicating, again, assignment to the Italian linguistic group) and the other covariates. The fitted values $\hat{F}$ and $F_i \times P_t$ from first stage regressions (6.2.1) and (6.2.2) are then plugged into equation (6.2.3) to predict populist preferences Y. These preferences are computed by multiplying the number of votes each party received in each municipality and for each election by the party and election-specific populist scores we obtained earlier.

What we develop is essentially a diff-in-diff model where treatment – exposition to fake news – is predicted by the assignment to the language group. In this way, the relationship

between linguistic groups and exposure to fake news can be tested in the first stages of our model (6.2.1) and (6.2.2), rather than naively assuming that the German-speaking population is completely unexposed to fake news as in our initial diff-in-diff framework, as we believe that membership to a linguistic community influences the intensity of exposition rather than fully determining it. If randomisation is achieved through assignment in a linguistic community and if we are able to control for individual qualities and common trends across the two groups, then the coefficient λ of the instrumented interaction term between predicted exposition and year of election will capture the causal effect of fake news exposure on electoral behaviour. An interesting implication of this new setting is that equation (6.1) will stand as the reduced form of this model, explaining the direct effect between the instrument and our outcome of interest.

7 Results

Table 2 shows the results from an initial linear model, using, as for all subsequent specifications, the natural logarithm of total populist scores by municipality, computed by aggregating text analysis scores obtained through the 'Assertive' dictionary, as the outcome variable. Estimates from columns (1) to (4) indicate that the positive correlation between exposure to fake news and populist vote is robust to the addition of various controls.

Column (1) presents our baseline model, including exposure to fake news, year of election (a dummy coded 1 for 2018, and 0 for 2013), their interaction, and the number of Italian-speaking voters by municipality. As expected, the greatest increases in populist vote are explained by the year dummy, coinciding with the rise of populist platforms in 2018. Nonetheless, the interaction between year and exposition, indicates a positive and significant effect of misinformation on voting in 2018.

Column (2) adds the number of landline low-latency connections to the covariate group. The introduction of the variable in the model increases the model precision from a 37.4% to a 50.6% r-squared statistic. Still, the exposition effect, given again by the interaction term, remains stable, and surpasses the coefficients for language group and internet connectivity in magnitude. The interaction coefficient is again unaffected from the introduction of controls for income per capita in Column (3).

Further controls for electorate size and abstention behaviour, increasing model fit to 97.9%, are added in Column (4). The negative link between abstention and populist vote is not unexpected and indicates that populist platforms may have attracted a number of individuals disenchanted by traditional politics. Most importantly, the log of the electoral size reveal itself as the most significant predictors in the model, indicating that the presence of a non-linear relationship between population density and populist vote. As a result, the effect of exposition on voting is reduced, but remains statistically significant, indicating that, for each additional

Table 2: OLS estimates of the effect of misinformation on populist vote

VARIABLES	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Exposed to fake news	-0.0011354***	-0.0012582***	-0.0012231***	-0.0001448***
	(0.0001642)	(0.0001845)	(0.0001742)	(0.0000305)
Exposed to fake news × Year of election	0.0007213***	0.0007971***	0.0007794***	0.0001188***
	(0.0001111)	(0.0001502)	(0.0001358)	(0.0000200)
Year of election	0.2594730***	0.1984275***	0.0584880	0.2994069***
	(0.0664220)	(0.0584740)	(0.0558583)	(0.0137647)
Italian speaking voters	0.0002711***	0.0000789**	0.0001118***	0.0000234***
	(0.0000294)	(0.0000354)	(0.0000359)	(0.0000057)
Broadband connections		0.0001959***	0.0001526***	0.0000423***
		(0.0000369)	(0.0000345)	(0.0000073)
Income per capita (natural log)			1.8080438***	-0.0444949
			(0.2453946)	(0.0602191)
Electorate size (natural log)				0.9826074***
				(0.0116515)
Abstentions and invalid votes				-0.0001800***
				(0.0000289)
Observations	584	584	584	584
Adjusted R-squared	0.374	0.506	0.560	0.979

*N*otes: Robust standard errors in parentheses. OLS estimates for the effect of misinformation on populist vote (natural log). Populist scores computed using the 'Assertive' text bag.

*p<.05; **p<.01; ***p<.001

like to a misinformation disseminator in a given municipality, populist scores increases by around 0.00011%. While this might seem like a minor effect, it should be noted that the interaction between exposure to fake news and the 2018 elections is still a significant predictor of populist vote whose order of magnitude, apart from electorate size and year of election, is still comparable to the effect of other covariates.

While showing a correlation between populist preference and exposition to misinformation, the results from our OLS specification may still, however, suffer from bias, as those estimates give no information on the direction of the causal channel. Indeed, we do not know whether access to misinformation bubbles is linked to individual characteristics that may already determine a populist preference, either through self-selection or online recruitment. These endogeneity issues cannot be addressed by simple correlations, and motivate our instrumental variable design.

Results from our two-stage-least-squares model (column 4) are shown in Table 3, including reduced form (1) and first stage (2 and 3) estimates. Here, the number of italian-speaking voters by municipality and its interaction with year of election are instrumented to predict exposition to fake news and its interaction with year. With the exception of the instrument and the endogenous exposition variable, the specifications in Table 3 retain the same control

group used in our final OLS model, as in column (4) of Table 2.

Reduced form estimates are displayed in column (1) of Table 3. As discussed earlier, this specification corresponds to the difference in difference model discussed in equation (1), where the effect of the exogenous variable on the outcome is given by the interaction term between language group shares and year of election.

Diff-in-diff estimates indicate a positive and statistically significant effect of the Italian language group on voting in the 2018 election. Accordingly, after controlling for all observables, the populist content of the vote in each municipality increased by 0.00001% for each additional Italian-speaking voter, during the 2018 general elections. These figures indicate a much smaller coefficient compared to the one of fake news exposition in the final OLS specification (Table 2, column 4). Reduced form figures are proportional to the causal effect of interest, therefore the positive sign might still indicate the presence of an effect. Nonetheless, unless we assume that exposition to misinformation is completely proxied by language groups, the coefficient of the interaction term is not a substitute for the effect of exposition to fake news on voting.

Turning at our instrumental variable estimates, first stage regressions predicting exposition to fake news and its interaction with year are shown in columns (2) and (3). The positive effect of the interaction between language group and year of election confirms our assumptions on differential exposure to fake news based on linguistic grouping. Results from column (3) indicate that, for each additional Italian-speaking voter, the number of likes to fake news disseminators increased by 0.44. Both models predict exposition with relative precision, boasting, overall, a 94.1% and a 93.7% adjusted r-squared. Partial r-squared statistics also indicate that, in both cases, most of the variation unexplained by the control covariates is captured by the instrument. Most importantly, both instruments pass the F-test for excluded instruments (as suggested in Bound et al., 1995), increasing our confidence in our estimates.

Our final estimates are presented in column (4), showing that fake news had no statistically significant effect on contributing to the rise of populist platforms. Indeed, the interaction coefficient between fake news exposition and year is not only lower than the OLS model in Table 2, but also not statistically different from zero. These results take any credit for the success of populist parties in 2018 away from the spread of fake news, and support the hypothesis that voters self-select into misinformation bubbles and consume fake news because of their prior preference for populist platforms, and not the other way around.

The untestability of the exclusion restriction in instrumental variable estimation implies that we cannot ensure that language group affiliation affects voting through channels other than exposition to fake news and year of election, for which we control through the baseline diff-in-diff framework. In this context, the estimates in Table 4, where a different instrument is used, offer a good robustness check of our results.

Table 4 presents an alternative specification where the number of landline broadband connections is instead instrumented. This use of internet penetration is not new in the literature,

Table 3: 2SLS estimates of the effect of misinformation on populist vote

	(1)	(2)	(3)	(4)
	DiD	First Stages		2SLS
VARIABLES		Exposition	Interaction	
Italian speaking voters	-0.0000001	0.1260896***	-0.0180990	
	(0.0000042)	(0.0384103)	(0.0279361)	
Italian speaking voters × Year of election	0.0000155***	0.2822714***	0.4468548***	
	(0.0000043)	(0.0538177)	(0.0521837)	
Year of election	0.2947903***	-299.6037610***	-391.4488535***	0.3085528***
	(0.0138577)	(61.4832894)	(59.8511025)	(0.0137742)
Broadband connections	0.0000434***	0.0268573	-0.0068932	0.0000435***
	(0.0000074)	(0.0438686)	(0.0328954)	(0.0000081)
Electorate size (natural log)	1.0008170***	-329.7326104***	-261.1409531***	1.0105005***
	(0.0097172)	(54.2358958)	(53.3322464)	(0.0133441)
Income per capita (natural log)	-0.0633365	517.8909599***	467.8603218***	-0.0804037
	(0.0601505)	(153.9726952)	(146.1050566)	(0.0596232)
Abstentions and invalid votes	-0.0001975***	0.1511572	0.1866705	-0.0002041***
	(0.0000310)	(0.1454624)	(0.1310689)	(0.0000356)
Exposed to fake news				0.0000039
				(0.0000336)
Exposed to fake news × Year of election				0.0000322
				(0.0000218)
Observations	584	584	584	584
Adjusted R-squared	0.979	0.941	0.937	0.978
Partial R-squared		0.729	0.843	
F-Test		15.57	69.91	

*N*otes: Robust standard errors in parentheses. IV estimates (including reduced form - DiD - and first stages) for the effect of misinformation on populist vote (natural log). Populist scores computed using the 'Assertive' text bag. F-tests for excluded instruments for the individual instrument (voters in the Italian-speaking language group) and its interaction with year of election are reported as F-Test (exposition) and F-Test (interaction), respectively.

*p<.05; **p<.01; ***p<.001

Table 4: 2SLS estimates of the effect of misinformation on populist vote (alternate instrument)

	(1)	(2)	(3)	(4)
	DiD	First Stages		2SLS
VARIABLES		Exposition	Interaction	
Broadband connections	0.0000465***	0.1038088**	0.1210049**	
	(0.0000082)	(0.0505354)	(0.0568417)	
Broadband connections × Year of election	0.0000156***	0.2938336***	0.4682898***	
	(0.0000040)	(0.0598440)	(0.0665250)	
Year of election	0.2921396***	-354.2210447***	-479.8780512***	0.3890398***
	(0.0139390)	(78.9100472)	(86.3438822)	(0.1216419)
Italian speaking voters	0.0000047	0.2079858***	0.1095598*	-0.0002111
	(0.0000047)	(0.0519421)	(0.0594555)	(0.0003096)
Electorate size (natural log)	1.0031237***	-281.7179699***	-183.2676752***	1.2633062***
	(0.0094949)	(55.3052189)	(58.2664905)	(0.3251746)
Income per capita (natural log)	-0.0683033	425.9142075***	321.8152469*	-0.4203606
	(0.0605966)	(158.1631244)	(164.0089705)	(0.5083296)
Abstentions and invalid votes	-0.0002307***	-0.5304766**	-0.9159087***	-0.0002677
	(0.0000347)	(0.2504293)	(0.2747019)	(0.0002180)
Exposed to fake news				0.0015240
				(0.0019172)
Exposed to fake news × Year of election				-0.0009230
				(0.0011707)
Observations	584	584	584	584
Adjusted R-squared	0.979	0.934	0.920	0.798
Partial R-squared		0.648	0.817	
F-Test		12.14	24.89	

*N*otes: Robust standard errors in parentheses. IV estimates (including reduced form - DiD - and first stages) for the effect of misinformation on populist vote (natural log). Populist scores computed using the 'Assertive' text bag. F-tests for excluded instruments for the individual instrument (number of broadband connections) and its interaction with year of election are reported as F-Test (exposition) and F-Test (interaction), respectively.

*p<.05; **p<.01; ***p<.001

as it has been instrumented in a number of studies to predict participation in internet-mediated activities (see, for example, Billari et al., 2018).

Surely, in a cross-sectional setting, the exogeneity of internet connectivity and its efficacy as a randomiser might be debatable, especially in contexts where rural areas are more difficult to reach by internet infrastructures. Yet, its connection with fake news exposition is reasonable, if not obvious, and the staggered investments in connectivity in the region may make a case for changes in the number of low-latency connections between 2013 and 2018 to still pose as an exogenous source of variation.

Reduced form estimates in column (1) show impressive similarities between the interaction term coefficient for the two instruments. The first stages in columns (2) and (3) also confirm that internet connectivity generates a positive effect on exposure to misinformation. According to column (3), the following of social media disseminators in 2018 increased by 0.46 likes for each household connected to a broadband connection. The precision of this alternative instrument is, however, reduced when compared to the previous one, with partial R-squared statistics indicating that the broadband instrument and its interaction with time only explains 64.8% and 81.7% of the remaining variation in exposition to misinformation. Yet, our concerns for instrument weakness are quickly dispelled, as both the instrument and its interaction with year manage to pass the F-tests for excluded instruments.

Column (4) shows the IV estimates and, while an increase in variability should be noted, the effect of exposition to fake news in 2018 is again not statistically different from zero. We then conclude that our year of election controls are good enough for our language group instrument not to affect our estimates, and that fake news has had no effect on voting outcomes in Trentino-Alto Adige/Südtirol over the course of the 2018 elections.

Table 5 shows estimates under a different outcome variable. In this figure, we leave populist scores aside and estimate the direct effect of misinformation on the electoral outcomes of the two major anti-establishment party platforms in the region. This model retains the same specifications from Tables 3 and 4, but focuses on the empirical connections – discussed in section 3.2 – between misinformation and the platforms of LEGA and M5S, testing whether these parties have benefited from any electoral gain by the introduction of fake news.

The first stage is identical to the ones reported in Tables 3 and 4, so they have been omitted, and the table only reports reduced forms and second stages under the two alternative instruments. Reduced form models, again, produce comparable results, revealing that each additional Italian-speaking voter increased vote to LEGA and M5S by 0.12 votes, while each additional access to low-latency connection increased this figure by 0.13.

As in our earlier models, both second stages (columns 2 and 4) display a smaller effect than what the reduced form coefficients suggest. The sign of this effect is, however, different, depending on the instrument used: while the broadband penetration instrument still suggests a non-significant effect on voting, the language group instrument indicates a -0.56 reduction

Table 5: 2SLS estimates of the effect of misinformation on vote to LEGA and M5S (binary outcome)

	(1)	(2)	(3)	(4)
	Language group		IT connectivity	
VARIABLES	DiD	2SLS	DiD	2SLS
Italian speaking voters	0.17663***		0.20725***	0.03842
	(0.00838)		(0.01061)	(0.27479)
Italian speaking voters × Year of election	0.11989***			
	(0.00928)			
Broadband connections	0.01619	-0.02315	0.06160***	
	(0.01265)	(0.06440)	(0.01656)	
Broadband connections × Year of election			0.13135***	
			(0.00804)	
Year of election	73.24935***	247.33200***	45.93550***	233.22003**
	(13.78798)	(46.84445)	(12.43112)	(111.98957)
Electorate size (natural log)	61.15233***	348.66490***	85.44525***	302.21239
	(12.16195)	(86.66988)	(11.66156)	(293.99039)
Income per capita (natural log)	10.53376	-408.42314***	-29.45080	-341.87453
	(37.38124)	(118.24482)	(36.52034)	(445.24376)
Abstentions and invalid votes	0.05411	-0.03984	-0.28461***	-0.07156
	(0.04848)	(0.18060)	(0.05609)	(0.21393)
Exposed to fake news		1.31970***		0.99183
		(0.35714)		(1.74009)
Exposed to fake news × Year of election		-0.56535**		-0.34184
		(0.25948)		(1.07813)
Observations	584	584	584	584
Adjusted R-squared	0.994	0.910	0.995	0.932

*N*otes: Robust standard errors in parentheses. IV estimates for the effect of misinformation on vote to LEGA and M5S (total votes by municipalities).

*p<.05; **p<.01; ***p<.001

for each additional like to a disseminator.

These differences are easily explained. As we will discuss later, this negative effect can be expected, as we will argue that the introduction of misinformation in a system might generate detrimental effects under specific conditions. However, the results from our internet connectivity instrument indicate that this might not be the case.

The lack of a continuous indicator oversimplifies the nuanced position of the SVP party which, as we discussed, still possesses the qualities of a catch-all party while still placing itself in the pro-establishment camp. There is a populist dimension to the SVP which only a continuous outcome can capture, meaning that a draconian assignment of the SVP into the 'non-populist' field will affect our estimates under the language group instrument when vote for this party specifically is so clearly correlated with the share of German-speaking voters.

This bias, however, is mitigated through the internet connectivity instrument, which is unaffected by language group allegiance. This non-significance of the interaction effect confirms our earlier results on the effect of fake news on voting, suggesting that the anti-establishment bias of misinformation in social media did not translate into a direct electoral advantage for these two parties.

Finally, alternative estimates, using the auxiliary populist score based on the 'Anti- establishment/ aggressive' text bag (as exposed in Section 5), are presented in Appendix B. The three tables replicate the ones presented in this section, with Table 7 showing OLS estimates and Tables 8 and 9 displaying IV estimates using language group and internet connectivity instruments, respectively. Our results and interpretations remain mostly unchanged. Table 7 shows that a positive correlation between populist preference and exposition is still present. Again, these results are overturned with our IV estimates from Tables 8 and 9, showing that exposition had a negative yet not significant effect on voting when treatment is assigned through random variations in language groups and internet connectivity.[35]

8 Discussions and conclusions

8.1 A simple model for misinformation and policy preferences

In this section we propose a simple model for misinformation and voting that could assist us in contextualising our results.

Our model draws on the median voter theorem, an approach not dissimilar to the one followed by Madestam et al. (2013). While simple, it can provide some useful conceptualisations which could help better understand the relationship between misinformation and voting preferences.

[35]First stages are reproduced for clarity purposes but, for obvious reasons, they are identical to the ones shown in Tables 3 and 4.

Suppose that two party platforms are contending for election, and that the position of these platforms can be drawn along a continuous policy axis representing support for the current establishment. The 'incumbent' platform sits on the left side of this axis, holding a pro-establishment stance, while the 'challenger' platforms sits on the opposite direction with an anti-establishment/populist stance.

Suppose each voter i has a prior preference g_i mirroring the same pro/anti establishment continuum characterising the two platforms. We assume these preferences are normally distributed with mean $\overline{g}$ and standard deviation σ.

Suppose also that each voter possesses an intrinsic ability a_i to ascertain the veracity of a given piece of information. This ability can follow any symmetric distribution as long as it is orthogonal to g_i. For clarity purposes, we will assume a_i follows a beta distribution with equal shape parameters, so that $A \sim Beta(\alpha, \beta)$ will be approximately normal while still being supported on the interval $[0, 1]$. a_i will then denote the percentage of fabricated facts which voter i will misreport as truthful, with $a_i = 0$ designating perfect ability to recognise a fake fact and $a_i = 1$ indicating that the voter will believe any piece of information he or she is presented with.

Vicinity determines voting preference. Given the two party platforms setting, there will be a level of support for policy g_{pop}, after which a voter will prefer to vote for the anti-establishment platform. Since g_i is normally distributed, the median voter 'bliss' preference will correspond to the mean of the population preference, meaning that, if the mean $\overline{g}$ equals g_{pop}, then the challenger will win the elections.

Now assume each voter is exposed to the same information stream, and receives a number of F facts. Each fact $j : [1, 2, ..., F]$ can either offer support for the anti-establishment movement (and be coded as $f_j = 1$) or show evidence against it (and be coded as $f_j = 0$). Information affects final support for the policy conditional on prior beliefs. We model final preference p_{i0} as a function of prior beliefs and information held by each individual $f(g_i, \sum_{j=1}^{F} f_{ji})$. We make the simple assumption that p_{i0} is directly proportional to g_i, given the share of facts in support of the anti-establishment, meaning that, if the information stream is the same for everyone, $p_{i0} = g_i(\sum_{j=1}^{F} f_j)/F$. Since a_i is bounded to the $[0, 1]$ interval, p_{i0} cannot be bigger than g_i, meaning that post-information distribution of policy preferences will have smaller mean and variance of the distribution of g_i. In any case, the normal distributional form of final preferences is unaffected, and the mean estimator will again indicate which platform the median elector will vote for. No further influx of misinformation will achieve support to the challenger platform if g_{pop} is already larger than the median value for g_i.

Now, suppose that a 'disseminator' is able to spread misinformation in support of the anti-establishment platform. We already discussed the empirical links between populism and misinformation, and we believe these factual connections can be justified by a number of factors. Indeed, it is possible for the pro-establishment platform to suffer from severe penal-

ties from disseminating blatantly false facts in its favour: due to its position as an institutional force, trust in the institutions might stand as a distinctive signifier of the political identity and communication of the platform. As a consequence, said platform might be particularly sensitive to erosion of public trust in traditional media and institutions, rendering a disinformation strategy potentially counter-productive. The pro-establishment platform, due to its incumbent position, may also have better access to traditional media, making the use of new communication technologies the only sensible choice for the challenger. In this case, the use of 'alternative facts' as a source of political legitimacy may devalue the role of traditional media sources, which instead stand as a vital asset for a pro-establishment platform.

In any case, we assume platforms are faced with a different set of incentives that make the dissemination of misinformation a less than optimal strategy for the incumbent platform, meaning that all misinformation will be in support for the challenger. We leave any further consideration on the nature of these asymmetries to future research, resting, for our intents and purposes, on the factual evidence supporting the empirical link between anti-establishment movements and fake news, as discussed in Section 3.2.

The disseminator is presented with a simple problem, to shift the median voter to g_{pop} by introducing misinformation. The disseminator has no control over who is going to be exposed to its fake facts, but it can affect the number of anti-establishment facts in the system. The insertion of these facts will increase the total stock of facts F by K, and will affect the individual share of perceived supported facts conditional on a_i. The formula yielding $p_{i,1}$, the misinformation-adjusted final preference, is given by equation 8.1.

$$\textbf{(8.1)}\quad p_{i,1} = \sum_{i=1}^{N} g_i \frac{\sum_{j=1}^{F} f_j + a_i K}{N(F+K)}$$

Gullible voters – for a_i approaching 1 – will take a share a_i of the total false facts K as true, and add it to their stock of perceived anti-establishment facts. Note that the introduction of fake news might generate a penalty: as the total stock of facts in the denominator is still increased by K, less gullible voters see their final misinformation-adjusted preference $p_{i,1}$ reduced when compared to $p_{i,0}$.

For large quantities of voters, the disseminator can replace a_i in equation 8.1 with its expectation $E[a_i]$. An unbiased estimator for $E[a_i]$ would then be given by the mean of a_i, as in equation 8.2. Given that, by design, $g_i \perp a_i$ and the distribution of a_i is symmetric, $p_{i,1}$ will still be symmetrically distributed, and its mean $\overline{p}_1$ will still reveal median voter preference $\tilde{p}_1$. An implication of this model is that $p_{i,1}$ and a_i are not orthogonal, so that, this time, support for the anti-establishment will positively correlate with ability to recognise truthful facts.

$$\textbf{(8.2)}\quad \tilde{p}_1 \mathrel{\hat{=}} \overline{g} \frac{\sum_{j=1}^{F} f_j + \overline{a} K}{F+K}$$

The misinformation-adjusted distribution of p_{i1} will then have mean as in the right side of equation 8.2 and standard deviation in $\sigma(\sum_{j=1}^{F} f_j + \overline{a}K)/(F+K)$. As $\overline{g}(\sum_{j=1}^{F} f_j + \overline{a})/(F+1)$

will denote the marginal effect of a single piece of misinformation being introduced, the number of fake facts to be introduced in order to shift the median voter to the anti-establishment platform will be given, after equating the right side of equation 8.2 to g_{pop} and rearranging, by:

(8.3) $K \hat{=} \frac{\overline{g}\sum_{j=1}^{F} f_j - g_{pop}F}{g_{pop} - \overline{g}\overline{a}}$

The disseminator will then only need to supply K fake news into the system for the anti-establishment platform to win the elections. In the context of linguistic groups affecting exposition, the subscript $l : [0, L]$ could be added, with K_1 set to zero[36] for voters belonging to linguistic minorities. In that case, when only $p_{i,1}$ is observed for at least two time periods, differences in the causal effect of exposition will be given by the difference between the change in final preference before $(t-1)$ and after (t) the introduction of misinformation across the two linguistic groups, as discussed in section (6): $\lambda = (\overline{p}_{1,t,0} - \overline{p}_{1,t-1,0}) - (\overline{p}_{1,t,L} - \overline{p}_{1,t-1,L})$.

While factual evidence still suggests that fake news have been overwhelmingly characterised by anti-establishment bias, our results, however, suggest that the electoral behaviour envisaged under this mechanism has not taken place during the Italian general elections in Trentino-Alto Adige/Südtirol.

There are two possible explanations for this outcome. First of all, mean ability might be too low. Should the distribution of A be skewed towards zero, or bounded by a support $[0, < 1]$, then lower values for $\overline{a}$ might lead to $\overline{p}_{i,1} < \overline{p}_{i,0}$, indicating that increases in support for the anti-establishment platform will be obtained for negative values of K. While factual evidence from the Italian elections suggests that $K > 0$, it could be argued that the anti-establishment platform possessed incomplete information on $\overline{a}$, overestimating the parameter. In these cases, a negative – or even null – effect of misinformation on final preferences is certainly possible. The preference-indifferent value for $\overline{a}^*$ can be obtained by solving for the first order condition of the partial derivative of $\overline{p}_1$ with respect to K in equation (8.2).[37] For values of mean ability lower than $\overline{a}^*$, any injection of misinformation will have detrimental effects on final preferences.

Our second explanation rests on a reasoning which transcends the model developed so far. Drawing on the echo chamber and filter bubble theories from Sunstein (2018) and Pariser (2011), it could be argued that a different mechanism should be introduced in order to explain the increased support to populist/anti-establishment platforms.

Our assumptions on information exposition may, indeed, be too unrealistic. There are no reasons to believe that every individual is exposed to the same stock of true facts as the next

[36] For simplicity, exposition is here set to zero for the least prevalent language group. Of course, degrees of exposition above zero are possible and have been accounted for in our regression model.

[37] Meaning that solving $(K+F)^2 = -\overline{g}(-\overline{a}F + \sum_{j=1}^{F} f_j)$ for $\overline{a}$ will yield the average ability level needed for fake news to have no effect. Preference-indifferent values for the share of true facts supporting the anti-establishment, holding ability and the other determinants of preference as fixed, can be similarly computed by solving this equation for $F_P = \sum_{j=1}^{F} f_j$.

one. It could be argued that each voter is exposed only to a fraction of information $N < F$, and that the share of true facts in support of a given platform is functional to the set of prior preferences of the voter. If $(\sum_{j=1}^{N} f_{ji})/N = f(g_i)$, then each voter is effectively exposed to the information he or she wants to believe, and the two platforms will play an entirely different game.

While modelling such a mechanism goes well beyond the scope of our study, the implications of this reasoning are almost obvious. In our model, already, as long as $\overline{a}$ is smaller than 1, then the marginal effect of an additional true fact in support of the anti-establishment will always generate greater electoral gains than an additional fabricated fact. Platforms may then have no incentive to disseminate misinformation if the stream of true information can already be tailored to the elector. In this context, fake news would arise because of a demand for facts supporting partisan views of the world, as an entirely natural process arising from the increased fragmentation and segregation of political opinion caused by the personalisation of social medial filter bubbles. In this case, the relationship between populist platform and the disseminator needs not to be symbiotic, and could as well be parasitic. The disseminator, while still creating false facts in support of the anti-establishment platform, may be driven entirely by economic motives – such as obtaining advertising revenue or increasing the value of his/her domain or fan page – with little to no voting externalities arising from his/her activities.

8.2 Conclusions

The measurement of the influence of fake news on electoral behaviour has, so far, escaped empirical assessment. With this study, we set to provide an answer to such a question, evaluating both the presence of this influence and, in that case, its magnitude.

In order to account for the inevitable reverse causality issues between voting preferences and exposition to misinformation, we proceeded with a quasi-experimental approach. Gathering municipality-level data on electoral outcomes, demographics and social media usage from the autonomous provinces of Trentino and South Tyrol, we exploited a natural experiment occurring in the region to randomise exposition to fake news.

In our contribution to the literature, we believe that our work sheds more light on the relationship between the spread of fake news and populist echo chambers. Mining the text from social media posts of parties and their leaders during their electoral campaign, we produced an indicator for populist content, allowing us to study populism as a phenomenon that eschews the political dimensions of left and right.

We showed that misinformation had a negligible and non-significant effect on the populist vote in Trentino and South Tyrol during the Italian 2018 general elections. Our results indicate that exposure to fake news is entirely dictated by self-selection in misinformation bubbles,

meaning that the causal channel between voting and fake news goes on a single direction, with individuals being exposed to misinformation because of their political presences.

In a simple two party model, if all voters were exposed to the same pieces of information, and if each voter had varying ability in recognising true facts from false ones, then misinformation could play a role in shifting median voter preferences. However, when social media filter bubbles are able to produce personalised information feeds, it is difficult to believe any given voter is exposed to the same information as the next one. Our results provide empirical evidence indicating either that the average ability to recognise true from false facts is underestimated, or that preferences are not dictated by the absolute proportion of facts perceived as true. This interpretation would support the hypothesis that the social media information bubbles can already shift the perceived proportion of supporting facts independently of the presence of misinformation.

This does not mean that the fake news is less problematic, but only that the causes of the populist shift in voting have to be found elsewhere, and that misinformation still thrives within these filter bubbles. The persistence of very similar differences in voting behaviour conditional on linguistic grouping and broadband penetration indicates that echo chambers most likely had a role in determining final preferences. In this sense, fake news would rather stand as the embodiment of shared narrations within groups of voters which are further reinforced by confirmation bias and the increasing personalisation of social media echo chambers. It may be possible that, once entered a misinformation bubble, partisan opinions are being reinforced by the presence of fake news, which might then ensure continued support for partisan beliefs as reputable sources of information are progressively removed and discredited in favour of 'alternative facts'. In the presence of personalised filter bubbles, preference dictates facts and not the other way around; social media plays a role in as much as it provides individuals with the information they want to believe.

Our final notes address the validity of our results and suggest future pathways for research, in the hope that our work spurs further empirical and theoretical contributions on the role misinformation has on voting.

The exploitation of a natural experiment Trentino-Alto Adige/Südtirol imposes a constraint on the external validity of our results, as the relationship between misinformation and voting might differ in other regional and national contexts. A similar methodology to ours could be applied to different contexts and with different units of observation, as language groups and broadband penetration have proven as good predictors for access to fake news. Survey data might also shed more light on individual preferences and social media behaviours.

Our estimates are also robust to a single definition of populism. Due to the liquid nature of this phenomenon, it is possible for our figures to change under different conceptions of populism, such as its more ethno-nationalistic departures. Our methodology could then be replicated using different text-bags as shown in Appendix B. In this sense, further work is

certainly needed to address competition between party platforms sharing contiguous filter bubbles and investigate whether misinformation can affect voting preferences between populist platforms, favouring certain versions of populism over other ones.

Finally, as argued earlier, we believe that future research should focus on the relationship between misinformation and echo chambers. As filter bubble access is already determined by prior preferences and individual characteristics, the attention of researchers and policy makers should rather focus on the socio-economic determinants for access into misinformation bubbles. If changes in voting behaviour are unaffected by increases in exposure to fake news, then, we believe, the personalisation of information streams in social media is ultimately the most important and socially-poignant factor that should be addressed when studying misinformation and the rise of populism.

Supplementary material

Datasets and codes used in our estimation can be found in the online data archive, available at the url address: https://sites.google.com/site/michelecantarella1992/data-archive-by-paper.

References

Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 44(2):139–177.

Albertazzi, D. and McDonnell, D. (2008). *Twenty-First Century Populism. The Spectre of Western European Democracy*. Palgrave Macmillan.

Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2):211–236.

Allcott, H., Gentzkow, M., and Yu, C. (2019). Trends in the diffusion of misinformation on social media. *Research & Politics*, 6(2).

Astat (2011). Censimento della popolazione 2011: Determinazione della consistenza dei tre gruppi linguistici della provincia autonoma di bolzano-alto adige. Astat Informazioni.

Astat (2017). Annuario statistico della provincia di bolzano. Provincia Autonoma di Bolzano/Alto Adige, Istituto provinciale di statistica - ASTAT.

Autorità per le Garanzie nelle Comunicazioni (2018). News vs fake nel sistema dell'informazione, interim report nell'ambito dell'indagine conoscitiva di cui alla delibera n. 309/16/cons. Interim report.

Azzimonti, M. and Fernandes, M. (2018). Social media networks, fake news, and polarization. Working Paper 24462, National Bureau of Economic Research.

Billari, F. C., Giuntella, O., and Stella, L. (2018). Broadband internet, digital temptations, and sleep. *Journal of Economic Behavior & Organization*, 153:58–76.

Bischof, D. and Senninger, R. (2018). Simple politics for the people? complexity in campaign messages and political knowledge. *European Journal of Political Research*, 57(2):473–495.

Bound, J., Jaeger, D. A., and Baker, R. M. (1995). Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American Statistical Association*, 90(430):443–450.

Boutyline, A. and Willer, R. (2016). The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, 38(3):551–569.

Caiani, M. and Graziani, P. R. (2016). Varieties of populism: insights from the italian case. *Italian Political Science Review*, 46(2):243–267.

Campante, F., Durante, R., and Sobbrio, F. (2017). Politics 2.0: The Multifaceted Effect of Broadband Internet on Political Participation. *Journal of the European Economic Association*, 16(4):1094–1136.

Card, D. and Krueger, A. (1994). Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *American Economic Review*, 84(4):772–93.

de Vreese, C. H., Esser, F., Aalberg, T., Reinemann, C., and Stanyer, J. (2018). Populism as an expression of political communication content and style: A new perspective. *International Journal of Press/Politics*, 23(4):423–438.

Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., and Quattrociocchi, W. (2016). Echo chambers: Emotional contagion and group polarization on facebook. *Scientific reports*, 6:37825.

Durante, R., Pinotti, P., and Tesei, A. (2019). The political legacy of entertainment TV. *American Economic Review*, forthcoming.

Ebner, C. V. (2016). The long way to bilingualism: the peculiar case of multilingual south tyrol. *International Journal for 21st Century Education*, 3(2):25.

Galbraith, J. K. and Hale, J. T. (2008). State income inequality and presidential election turnout and outcomes. *Social Science Quarterly*, 89(4):887–901.

Giglietto, F., Iannelli, L., Rossi, L., Valeriani, A., Righetti, N., Carabini, F., Marino, G., Usai, S., and Zurovac, E. (2018). Mapping italian news media political coverage in the lead-up of 2018 general election. *SSRN Electronic Journal*.

Guess, A., Nyhan, B., and Reifler, J. (2018). Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 us presidential campaign. *European Research Council*.

Guiso, L., Herrera, H., Morelli, M., and Sonno, T. (2017). Populism: Demand and supply. *Center for Economic Policy Research Discussion Paper*, 11871.

Gunther, R., Beck, P. A., and Nisbet, E. C. (2019). "Fake news" and the defection of 2012 obama voters in the 2016 presidential election. *Electoral Studies*.

Hersh, E. D. and Nall, C. (2015). The primacy of race in the geography of income-based voting: New evidence from public voting records. *American Journal of Political Science*, 60(2):289–303.

Kaltwasser, C. R., Taggart, P. A., Espejo, P. O., and Ostiguy, P. (2017). *The Oxford Handbook of Populism*. Oxford University Press.

Kumar, A. and Sebastian, T. M. (2012). Sentiment analysis on twitter. *International Journal of Computer Science Issues*, 9(4):372–378.

Lewis-Beck, M. S. and Nadeau, R. (2011). Economic voting theory: Testing new dimensions. *Electoral Studies*, 30(2):288–294.

Madestam, A., Shoag, D., Veuger, S., and Yanagizawa-Drott, D. (2013). Do political protests matter? evidence from the tea party movement. *The Quarterly Journal of Economics*, 128(4):1633–1685.

Martin, G. J. and Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–2599.

Nyhan, B. and Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, 32(2):303–330.

Pariser, E. (2011). *The Filter Bubble*. Penguin Books Ltd.

Pauwels, T. (2011). Measuring populism: A quantitative text analysis of party literature in belgium. *Journal of Elections, Public Opinion and Parties*, 21(1):97–119.

Polk, J., Rovny, J., Bakker, R., Edwards, E., Hooghe, L., Jolly, S., Koedam, J., Kostelka, F., Marks, G., Schumacher, G., et al. (2017). Explaining the salience of anti-elitism and

reducing political corruption for political parties in europe with the 2014 chapel hill expert survey data. *Research & Politics*, 4(1):2053168016686915.

Rooduijn, M. (2014). The nucleus of populism: In search of the lowest common denominator. *Government and Opposition*, 49(4):572–598.

Rooduijn, M. and Pauwels, T. (2011). Measuring populism: Comparing two methods of content analysis. *West European Politics*, 34(6):1272–1283.

Roozenbeek, J. and van der Linden, S. (2018). The fake news game: actively inoculating against the risk of misinformation. *Journal of Risk Research*, pages 1–11.

Schkade, D., Sunstein, C. R., and Hastie, R. (2007). What happened on deliberation day. *California Law Review*, 95:915.

Shin, J., Jian, L., Driscoll, K., and Bar, F. (2018). The diffusion of misinformation on social media: Temporal pattern, message, and source. *Computers in Human Behavior*, 83:278–287.

Shu, K., Sliva, A., Wang, S., Tang, J., and Liu, H. (2017). Fake news detection on social media. *ACM SIGKDD Explorations Newsletter*, 19(1):22–36.

Sunstein, C. R. (2002). *Republic.com*. Princeton University Press.

Sunstein, C. R. (2018). *#Republic*. Princeton University Press.

Taggart, P. A. (2000). *Populism*. Open University Press.

Törnberg, P. (2018). Echo chambers and viral misinformation: Modeling fake news as complex contagion. *PLOS ONE*, 13(9):e0203958.

Appendix A: measuring exposition to fake news

The problem of measuring exposition to fake news is far from a trivial one. Other studies have relied on different methods, such as individual level survey data or the total social media shares for each piece of misinformation. The structure of our data, however, imposes to obtain a proxy for exposition that can be broken down at municipality level and that is available for all units of analysis. There is no way to reconstruct how each fake news has been shared across social media, and who has been exposed to it. It is, however, possible to construct an approximate measure of the following each Facebook page which disseminate fake news has in a given municipality.

This information can only be acquired through a single source: the Facebook Audience Insight Tools (see www.facebook.com/ads/manager/audiences). This tool, generally intended for advertising purposes, allows to access information on active Facebook users, who can be filtered by their place of residence, age, gender, language and many other interests. The tool yields demographic (broken down by age, gender, relationship status, education, and employment) and online activity (page likes and social media use) information on all active individuals in the targeted audience. While it is not possible to filter an audience by their appreciation for a particular Facebook page, it is possible to target an audience through 'interests', and then be presented with a list of pages which correlate with these interests amongst the selected audience, where the number of audience-specific likes for each of these pages is also displayed.

We found a number of interest to be correlated with appreciation of fake news disseminators in Italy. Keyworks such as 'Immigrazione', 'Stato sovrano', 'Scienza di confine', 'Illuminati', 'Medicina alternativa', 'Casta', 'Teoria del complotto', 'Notizie Incredibili', 'Popolo', 'News24', 'Sovranità', 'La casta', 'Massoneria', and 'Notizie.it', all returned information on many disseminators included in 'Black lists' compiled by debunking websites.

While the demographic information of the audience tool is biased towards Facebook users, we are confident that page likes which is information obtained from these tools presents an unbiased figure of exposition to misinformation, as it is expected for the the utmost majority of fake news to travel through social media.

We then used the Facebook Audience Insight Tools to extract a sample of pages likes on Facebook, using the 'News24' interest. We collected this information for each Municipality in the Trentino-Alto Adige/Südtirol region, for all users aged 24+, in order to ensure that these individuals had the opportunity to vote at both 2013 and 2018 elections.

The collection process required a few precautions to be taken. First, as the tool will not display figure on a location if the number of users in that area is below a certain (unspecified) threshold, we collected all exposition information for pairs of municipalities $Trento + Municipality_i$, using Trento (capital of the Trentino-Alto Adige/Südtirol) as the baseline to

be subtracted at the end of the collection. This method ensured that: (1) information could be collected for all locations, and (2) that the tool always returned the same list of pages.

A second issue is generated by the 'interest' targeting options, as the tool only returns information on the number of users belonging to a given 'interest' audience who also liked that page. This means that the number of observed likes through interest targeting will be a fraction – approaching unity, for certain interests – of the total audience, and that exposition figures will be under-estimated. To adjust for this issue, we make two further assumptions: (1) that the total audience of these pages is composed solely of people residing in Italy and (2) that this fraction is the same for both Trentino Alto Adige/Südtirol and Italy. More formally, we assume that $obslikes_{IT}/totlikes_{IT} = obslikes_{TA}/totlikes_{TA}$, meaning that the ratio between observed likes and total likes is equal across the country. Under this assumption, we divide, for each page, the total number of likes by the audience figures obtained in the Italian territory, and we adjust the shares collected in Trentino Alto Adige/Südtirol by this scalar.

Finally, the inability to target specific pages through the audience tool implies that not all fake news disseminators will be captured by this methodology, as some are – perhaps purposely – unreachable through interest targeting.

To account for this issue, we develop a simple yet powerful predictive model that we use to impute the social media following for each of these missing pages, following from the intuition that the number of individuals liking a page in a given municipality will be proportional to the total Facebook likes of said page, holding specific effects from each location, such as its size and the language group, as fixed. If we assume that all page likes follow a similar functional form, we can then estimate the parameters from this function for observed pages, and then use the model to impute the following of unobserved pages. In the model:

$$y_i = \alpha + totlikes_i\beta + municip_i'\gamma + totlikes_i \times municip_i'\delta + collegio_i'\zeta + e_i$$

y_i indicates the total number of Facebook likes each observed page has in each municipality, for $i = \{1, 2, ..., M \times P\}$, where M indicates the total number of municipalities, and P the total number of observed pages. $Totlikes$, instead, indicates the total number of likes each of these pages has on Facebook, while $Municip$ and $Collegio$ are column vectors of dummies for municipality and constituency, respectively. The pages used in our estimation are presented in Table 6, where observed and modelled pages are labelled as 'donor' and 'recipient' respectively.[38]

The model allows to predict the number of likes for an observed page in a given location as a function of the total number of likes of said page, allowing for different slopes and intercepts

[38] A very similar imputation model was also used to correct figures for the 'Chedonna.it' page. As the exposition figures are rounded by the nearest hundred after a page reaches 1,000 likes in a given location, figures for this have been adjusted accordingly using the variation in the other observed pages.

between municipalities.[39] With 3504 observations and a multiple R-squared of 0.9237, the model achieves a satisfying fit and its parameters are extracted and used to impute, for each municipality, the number of people liking the other pages appearing in 'black lists' using the total number of likes of these pages only. After these values have been imputed, they are summed with the observed values in each location to obtain an estimate of the total number of social media likes to fake news disseminators in each municipality.

While it is reasonable to argue that the fake news phenomenon has risen to mainstream attention only in the eve of the 2017 referendum and 2018 elections, it should be noted that few of these pages already existed before 2013. To account for this issue, we adopt a 'conservative' methodology and, for 2013, exposition to fake news is still computed using pages which existed in the date of the 2013 elections, keeping the total number of likes for each of these pages unaffected. While the results shown in Section 7 make use of these adjusted figures, our estimates are nearly unaffected by the use of a less conservative indicator where exposition to fake news is set at zero for 2013.

As a final note, it should be remarked that this imputation method is not stochastic, but deterministic. However, as the final variable will consist in the sum of these estimates, we are generally uninterested in correctly simulating within-municipality variation, while we feel that between-municipality variation is a second-order problem considering the good fit of the model. Also, we remark that the purpose of this imputation is to improve the figures of exposition so that the effect of each additional like to a disinformation disseminator can be quantified with more precision. In any case, even in the presence of over or under-estimation of our exposition figures, the sign, and the statistical presence of an effect of fake news on voting should remain unchanged, as the imputation mostly scales the total number of likes in a municipality upwards. Indeed, we constructed another alternative variable for exposition to misinformation using only information from the pages we managed to observe with the Facebook Audience Insight Tool, and our final figures are again unaffected by these changes.

Last but not least, our results are only partially reproducible, as the Facebook API suffers from severe transparency issues that affect the possibility to perfectly replicate this method. It is well possible, then, that exposition figures may vary slightly when collected in a different moment in time. Also, some pages may have changed their name or have been unpublished in the meanwhile, complicating, again, the reproduction of our results. Indeed, many of the pages reported in Table 6 were closed by Facebook in May 2019 following a flagging campaing from the non-governmental organisation AVAAZ.[40] In any case, the original data set used in our calculations is made available in the online data archive.

[39] The coefficient vector γ will produce fixed effects specific to the municipality, while the vector of the interaction coefficients δ will allow for different slopes.

[40] Ibid. 18

Table 6: Fake news disseminators

Page	likes	followers	founded in	type	Page	likes	followers	founded in	type
Adesso Basta	508971	522664	06-Apr-16	donor	Notizie live	11309	11246	03-Nov-14	recipient
Chedonna.it	1996969	1905609	03-Dec-13	donor	Questa è l'Italia	3982	3987	26-Jan-12	recipient
Citazioni che ispirano	322977	324125	02-Nov-13	donor	Questa è l'Italia di oggi	19674	20961	01-Aug-16	recipient
DiariodelWeb.it	506535	512444	19-Dec-14	donor	Rimani informato	44758	53156	03-Jul-17	recipient
Dimissioni e tutti a casa	784505	749631	19-Aug-13	donor	Rothschild: la Bestia che domina il mondo	19172	19163	21-Jul-11	recipient
Giornale Interattivo	525468	525468	03-Mar-14	donor	Scienza di confine	81788	81060	11-Feb-12	recipient
Italia Patria Mia	375094	369117	29-May-15	donor	Segreto di Stato	6918	6934	07-Dec-16	recipient
Mister Link	436253	404946	11-Apr-11	donor	Se ti fai un'altra birra la prima non s'incazza	205564	204537	11-Nov-12	recipient
Quello che i TG non dicono	369617	377979	23-Oct-16	donor	SocialTv Network (websocialtv)	99181	107729	05-Mar-16	recipient
Silenzi e falsità della stampa italiana	853552	861987	06-Nov-13	donor	Sono senza parole	938033	905287	29-Jun-14	recipient
24H Italia News	171561	173703	07-Dec-16	recipient	Sovranità Monetaria & Debito Pubblico	2128	2158	16-Nov-11	recipient
Aprite gli occhi	148648	149245	08-Feb-11	recipient	Stop alle scie chimiche	32729	32471	19-Oct-09	recipient
Autismo e Vaccini	22391	23016	22-Jul-13	recipient	Ultimissime	2889	2876	29-Apr-14	recipient
Avvistamenti di Creature Mitologiche	181711	179820	20-Mar-12	recipient	Un male tutto italiano, questo regime è una vera piaga sociale	10912	10686	10-Jun-11	recipient
Banda Bassotti	32666	32776	29-Jan-13	recipient	Vaccini Basta	27314	28314	21-Jan-12	recipient
Catena Umana	34464	34372	16-May-14	recipient	Zona grigia	66149	65974	29-Apr-16	recipient
Come i treni a vapore	132131	136164	07-Feb-14	recipient	Stop Invasione	18197	18718	23-Oct-15	recipient
Contro i poteri forti	11441	11591	23-May-15	recipient	Un caffe al giorno	235717	236577	21-Jul-16	recipient
Cose che nessuno ti dirà di nocensura.com	1666222	1612016	16-Aug-10	recipient	Non cielo dicono	153986	160074	25-Jun-16	recipient
CrimeNews (mafiacapitale.info) (Catena Umana 3)	11804	11813	13-Oct-15	recipient	Roby	1425198	1414607	06-Apr-11	recipient
CSSC - Cieli Senza Scie Chimiche	10531	10619	08-Dec-11	recipient	Sardegna Today	5428	5550	04-Sep-18	recipient
Eco(R)esistenza	85473	84418	02-Jun-11	recipient	Dangerous News	5385	5275	07-Jul-13	recipient
Informati (informatitalia)	211461	207818	07-Apr-13	recipient	Il Matto Quotidiano	12313	11952	02-Apr-15	recipient
IO VI Spengo	481267	550151	12-Apr-15	recipient	Imola Oggi	9188	9403	16-Sep-18	recipient
Italia Malata	96781	96935	02-Sep-15	recipient	Italiani Uniti per la Patria	96382	98380	28-Apr-16	recipient
Killuminati Soldiers	2406	2434	23-Jan-16	recipient	Scenari Economici	39224	42555	06-Mar-13	recipient
John Koenig Mb Koenig	2585	2670	16-Feb-14	recipient	Informare X Resistere	1086691	1020283	16-Jul-09	recipient
L'alimentazione e gli illuminati	10032	10095	10-Aug-11	recipient	Piovegovernoladro	55839	54235	19-Dec-13	recipient
L'antipolitica	14517	14519	29-Apr-12	recipient	Jeda News	372364	357423	02-Jan-12	recipient
La Vera Italia	3742	3722	16-May-15	recipient	Breaknotizie	81656	80762	10-Sep-12	recipient
La Verità ci Rende Liberi	68857	72846	20-Jun-18	recipient	Controinformo	27095	26917	13-Jan-16	recipient
Le notizie che non ti aspetti	133697	133185	26-Jun-14	recipient	DirettaNews24	130449	130453	09-Oct-16	recipient
Mafia Capitales (whatsappecultura)	99487	98732	30-Dec-08	recipient	FascinAzione	1866	1910	03-Apr-16	recipient
Movimento anti NWO	12324	12274	20-Feb-12	recipient	TG 5 Stelle	16934	16216	26-Jun-13	recipient
Neovitruvian	6742	6763	09-Dec-10	recipient	CVDiariodelPollineo	1569	1572	14-Oct-16	recipient
NO alla dittatoriale Unione Europea e al Nuovo Ordine Mondiale	2629	2714	11-Nov-11	recipient	Notizie in Movimento	116951	118686	29-Sep-14	recipient
Notiziario (notiziario.face)	22882	22943	15-Nov-12	recipient					

*N*otes: Donor and recipient pages used in the estimation of total page likes by municipality. Previous page names are shown in parentheses.

Appendix B: results under alternative populist score

Table 7: OLS estimates of the effect of misinformation on populist vote

VARIABLES	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Exposed to fake news	-0.0014947***	-0.0015837***	-0.0015547***	-0.0006203***
	(0.0002192)	(0.0002193)	(0.0002170)	(0.0001080)
Exposed to fake news × Year of election	0.0009528***	0.0010078***	0.0009931***	0.0004451***
	(0.0001458)	(0.0001610)	(0.0001541)	(0.0000683)
Year of election	0.1235975**	0.0793249	-0.0364128	0.2046793***
	(0.0613918)	(0.0571435)	(0.0555898)	(0.0298021)
Italian speaking voters	0.0003407***	0.0002013***	0.0002285***	0.0001351***
	(0.0000400)	(0.0000441)	(0.0000453)	(0.0000232)
Broadband connections		0.0001420***	0.0001062***	0.0000591***
		(0.0000318)	(0.0000300)	(0.0000186)
Income per capita (natural log)			1.4953522***	-0.1794182
			(0.2213559)	(0.1196278)
Electorate size (natural log)				0.8931439***
				(0.0238180)
Abstentions and invalid votes				-0.0003172***
				(0.0000736)
Observations	584	584	584	584
Adjusted R-squared	0.479	0.545	0.579	0.905

*N*otes: Robust standard errors in parentheses. OLS estimates for the effect of misinformation on populist vote (natural log). Populist scores computed using the 'Anti-establishment/aggressive' text bag.

*p<.05; **p<.01; ***p<.001

Table 8: 2SLS estimates of the effect of misinformation on populist vote

	(1)	(2)	(3)	(4)
	DiD	First Stages		2SLS
VARIABLES		Exposition	Interaction	
Italian speaking voters	0.0000459***	0.1260896***	-0.0180990	
	(0.0000173)	(0.0384103)	(0.0279361)	
Italian speaking voters × Year of election	0.0000268*	0.2822714***	0.4468548***	
	(0.0000155)	(0.0538177)	(0.0521837)	
Year of election	0.2148922***	-299.6037610***	-391.4488535***	0.2562892***
	(0.0337123)	(61.4832894)	(59.8511025)	(0.0365017)
Broadband connections	0.0000447*	0.0268573	-0.0068932	0.0000344
	(0.0000229)	(0.0438686)	(0.0328954)	(0.0000355)
Electorate size (natural log)	0.9827922***	-329.7326104***	-261.1409531***	1.0548130***
	(0.0213145)	(54.2358958)	(53.3322464)	(0.0406127)
Income per capita (natural log)	-0.2918652**	517.8909599***	467.8603218***	-0.3959852***
	(0.1316488)	(153.9726952)	(146.1050566)	(0.1394922)
Abstentions and invalid votes	-0.0003442***	0.1511572	0.1866705	-0.0003668***
	(0.0000898)	(0.1454624)	(0.1310689)	(0.0001194)
Exposed to fake news				0.0003419*
				(0.0001918)
Exposed to fake news × Year of election				-0.0001560
				(0.0001440)
Observations	584	584	584	584
Adjusted R-squared	0.8884455	0.9415209	0.9372169	0.8481310
Partial R-squared		0.729	0.843	
F-Test		15.57	69.91	

*N*otes: Robust standard errors in parentheses. IV estimates (including reduced form - DiD - and first stages) for the effect of misinformation on populist vote (natural log). Populist scores computed using the 'Anti-establishment/aggressive' text bag. F-tests for excluded instruments for the individual instrument (voters in the Italian-speaking language group) and its interaction with year of election are reported as F-Test (exposition) and F-Test (interaction), respectively.

*p<.05; **p<.01; ***p<.001

Table 9: 2SLS estimates of the effect of misinformation on populist vote (alternate instrument)

	(1)	(2)	(3)	(4)
	DiD	First Stages		2SLS
VARIABLES		Exposition	Interaction	
Broadband connections	0.0000632**	0.1038088**	0.1210049**	
	(0.0000250)	(0.0505354)	(0.0568417)	
Broadband connections × Year of election	0.0000337**	0.2938336***	0.4682898***	
	(0.0000154)	(0.0598440)	(0.0665250)	
Year of election	0.2060661***	-354.2210447***	-479.8780512***	0.3444824*
	(0.0344743)	(78.9100472)	(86.3438822)	(0.1834180)
Italian speaking voters	0.0000500***	0.2079858***	0.1095598*	-0.0002301
	(0.0000169)	(0.0519421)	(0.0594555)	(0.0004603)
Electorate size (natural log)	0.9907994***	-281.7179699***	-183.2676752***	1.3300134***
	(0.0203888)	(55.3052189)	(58.2664905)	(0.4918501)
Income per capita (natural log)	-0.3014202**	425.9142075***	321.8152469*	-0.7625836
	(0.1307629)	(158.1631244)	(164.0089705)	(0.7723294)
Abstentions and invalid votes	-0.0004524***	-0.5304766**	-0.9159087***	-0.0004730
	(0.0001085)	(0.2504293)	(0.2747019)	(0.0003403)
Exposed to fake news				0.0019554
				(0.0028663)
Exposed to fake news × Year of election				-0.0011550
				(0.0017511)
Observations	584	584	584	584
Adjusted R-squared	0.8898863	0.9342932	0.9202511	0.4898130
Partial R-squared		0.648	0.817	
F-Test		12.14	24.89	

*N*otes: Robust standard errors in parentheses. IV estimates (including reduced form - DiD - and first stages) for the effect of misinformation on populist vote (natural log). Populist scores computed using the 'Anti-establishment/aggressive' text bag. F-tests for excluded instruments for the individual instrument (number of broadband connections) and its interaction with year of election are reported as F-Test (exposition) and F-Test (interaction), respectively.

*p<.05; **p<.01; ***p<.001

Appendix C: correlation between text-based populist scores and CHES data

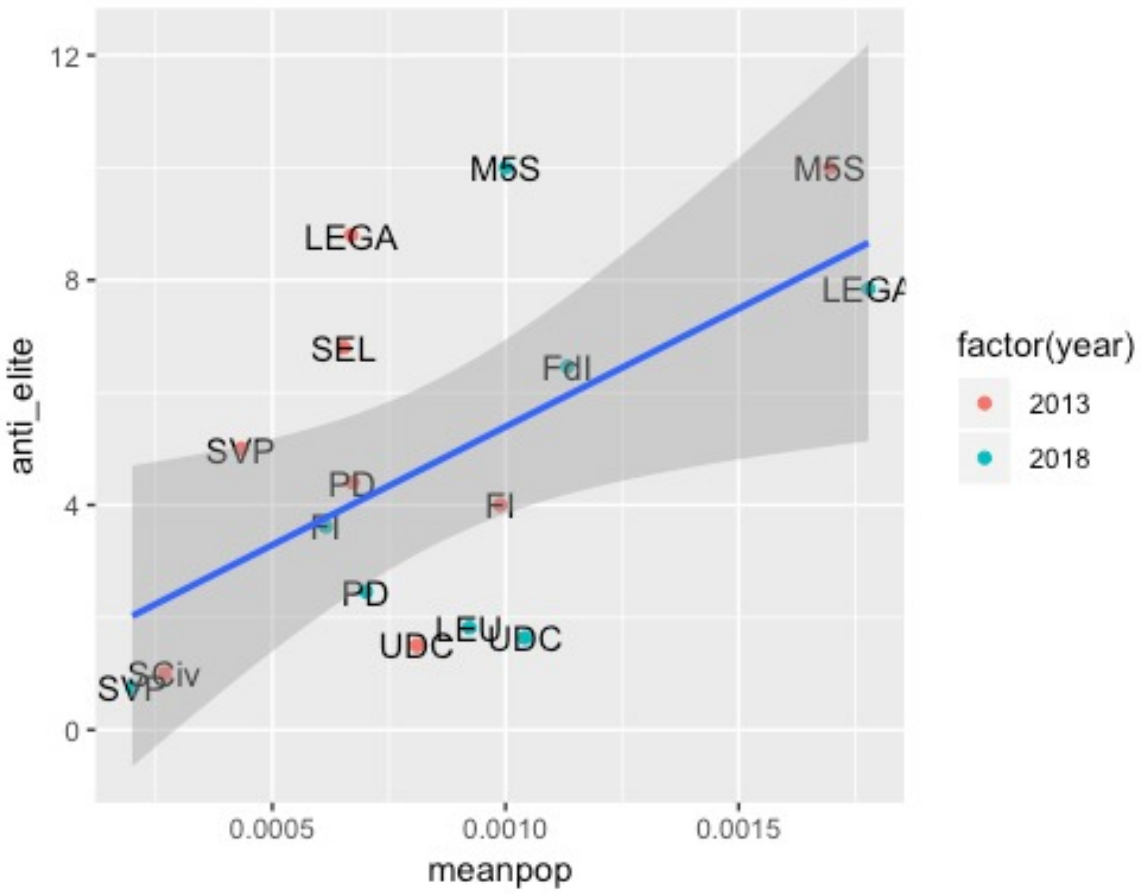

Figure 4: Text analysis scores of social media posts from parties and their leaders during the 2013 and 2018 elections campaigns (x-axis) and their relationship with the scores on the variable 'People vs the Elites' from the Chapel Hill Expert Survey for year 2014 (y-axis). Higher values on the y-axis correspond to higher salience of anti-establishment and anti-elite rhetoric on a scale from 0 to 10. Scores are computed using the 'Anti-establishment/ aggressive' text bag.

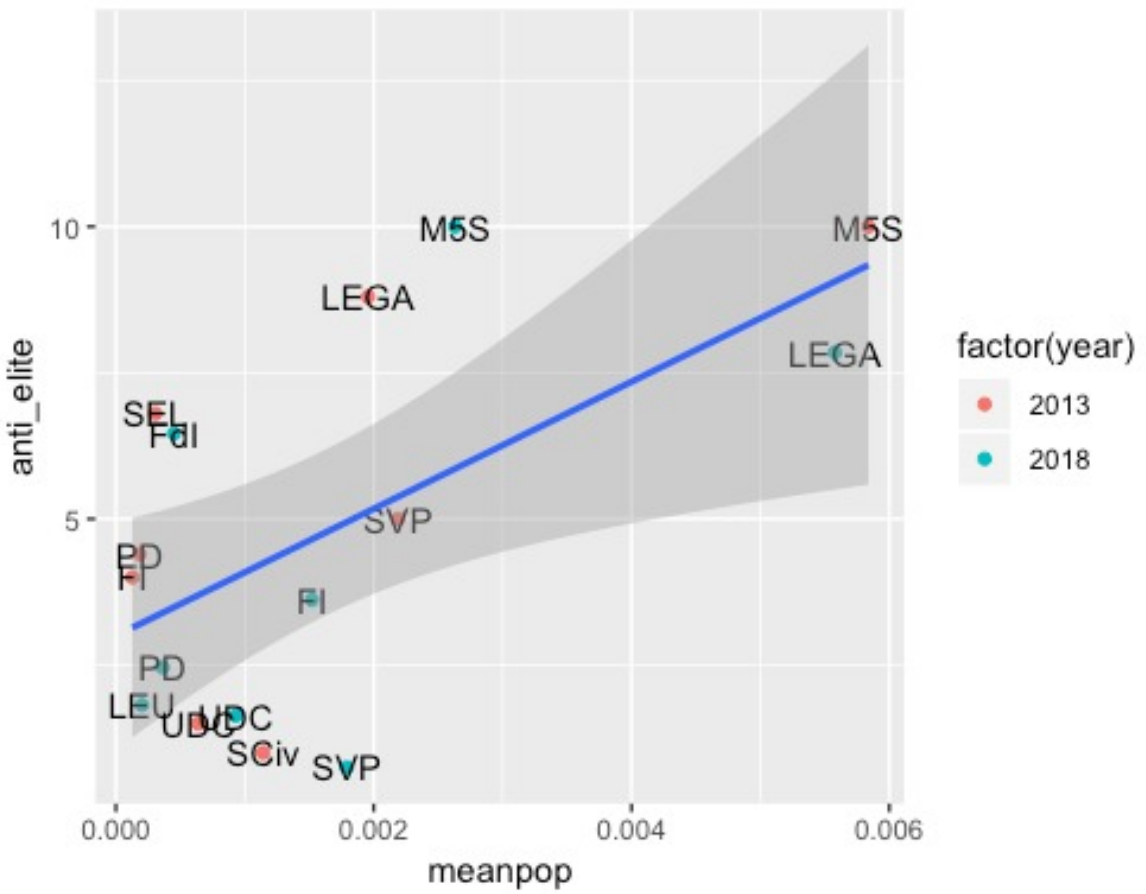

Figure 5: Text analysis scores of social media posts from parties and their leaders during the 2013 and 2018 elections campaigns (x-axis) and their relationship with the scores on the variable 'People vs the Elites' from the Chapel Hill Expert Survey for year 2014 (y-axis). Higher values on the y-axis correspond to higher salience of anti-establishment and anti-elite rhetoric on a scale from 0 to 10. Scores are computed using the 'Assertive' text bag.