Micro-Targeting and ICT media in the Dutch Parliamentary system

Technological changes in Dutch Democracy

Haye Hazenberg, Jeroen van den Hoven, Scott Cunningham, Mark Alfano, Hadi Asghari, Emily Sullivan, Amir Ebrahimi Fard and Elsa Turcios Rodriquez.

Table of Contents

Introduction

For the period surrounding the 2018 Dutch municipal elections, a team of researchers from the Delft University of Technology investigated the effect of the digital environment on parliamentary democracy. An interdisciplinary group of researchers combined expertise on digital ethics, political theory, big data analytics, the economics of privacy and security, epistemology, media studies and computer science. This report presents the main findings, which are grouped around two main themes: political micro-targeting and ICT media. Societal themes that came to prominence over the research period, such as the debate over 'fake news' and the leaks of personal information that were used for political purposes by Facebook, as well as the implementation of new EU privacy regulation helped to put the research in a larger political context.

The main findings provide a qualified picture. The influence of the digital revolution on democratic politics is already revolutionary, and the weaknesses of online platforms provide ample opportunities for derailing liberal democracy. Digital platforms are too closed-off, not mindful enough of individual digital rights, and biased in their (re)presentation of political pluralism. But the Netherlands has proven to be one of the few democracies that is relatively resilient, with an open multi-party system receptive to the political fragmentation that ICT developments encourage, and relatively high trust between citizens, in shared media organizations, and between political parties.

In order not to be complacent in the face of fundamental challenges, the report provides several urgent recommendations. Next to several 'reactive' recommendations, which seek to remedy the weaknesses and dangers the digital environment poses to democracy, it also outlines an example of how the digital environment might be proactively redesigned in order to positively enhance the quality of the Dutch parliamentary system.

This research was made possible by additional funding from the State Commission on the Parliamentary System. The authors are fully responsible for the contents of this publication. There has been no interference from the State Commission with this research and the conclusions and/or the recommendations in this publication do not necessarily reflect the opinions of the State Commission.

Chapter 1. Context and Design of the Research

Research Design

In the research proposal, a set of research questions were outlined to ascertain the impact of online information technologies on trust in both the recent Dutch municipal election process and the institutions of law-based representative parliamentary democracy more broadly. It was hypothesized that the immediate current threat to such trust concerned the de-legitimation of the democratic process itself as being fair, impartial, representative and open to deliberation through online 'echo chambers', or the tendency of online media consumption to push democratic citizens further down filter bubbles. The pluralism so essential to representative party politics can turn into democratic disinformation, when citizens are presented with information unchecked by the deliberative process that characterizes democracy.

The original research questions were as follows:

Main-RQ: Does ICT media de-legitimize the democratic municipal campaign and election process, and if so, to what extent and how?

Sub-RQ1: In which (echo chamber/filter bubble) directions does online political discourse develop during the process? (Twitter, forums/websites as possible)

Sub-RQ2: What particular tracking mechanisms (companies and websites) create an online advertising environment conducive to filter bubbles and echo chambers? (General browsing,)

Sub-RQ3: Which algorithms, if any, and which social bots, if any, steer online information consumption in one direction or the other? (Youtube, Twitter, websites)

Sub-RQ4: Which themes, political parties and politicians in the city of the Hague and the Netherlands more broadly characterize the municipal elections?

Sub-RQ5: How does Sub-RQ4 relate to Sub-RQ 1,2 and 3?

Of particular interest to democratic politics in the Dutch parliamentary system is the more focused mechanism of 'political micro-targeting'. Within the more 'macro' study of the effect of the ICT environment on digital echo chambers that can be of concern to democratic pluralism, the micro-action of political micro-targeting was highlighted as potentially most concerning. Political micro-targeting allows political actors to target individual citizens directly in their online browsing, so that every user gets a different, specifically tailored advertisement to see. This tailored advertising model is inherent to the economy of the current ICT environment, but could be potentially disastrous for democratic politics when political advertisements remain 'dark', or unknowable to anyone except the sender, receiver, and the digital platform intermediary. At that point, customized political advertisements could erode the public aspect of democratic politics. Different political actors should be able to contest one another's ideas in public, and the closed loop of

micro-targeted ads might prevent them from knowing about the ideas of their competitors. Micro-targeting might also mislead citizens into believing political actors publicly stand for something based on the ads they privately receive, and leave citizens unsure as to the priorities of political parties.(Zuiderveen Borgesius et al., 2018)

Further worries might arise from the potential threat to privacy and individual autonomy micro-targeting can pose. Online advertising technology has advanced far beyond the simple displaying of messages on pages, as it makes use of a wide array of both general psychological behavioural insights and very specific individual profiles to target advertisements at the exact right time and place and to the exact right person.(Helbing et al., 2017) Even as these ads might seem ineffective or overrated to most users, multi-billion dollar investments, and increasingly political careers, are staked on them being effective.

Next to this particular 'micro-targeting' mechanism, the larger customized advertisement-based ICT environment can also be potentially harmful for democracy in three further ways, which the sub-RQ's refer to. The rise of a 'new online public space' can be harmful when the 'old' shared public space is left vacant, or when it leads to further divisions. A contest of ideas between political actors can then seemingly take place in the 'old' public space, while citizens focus their attention instead on online information that is customized for them and fails to resemble the old public space in democratically meaningful ways. The ICT environment might for example not be very hospitable to nuanced debate or to ensuring that a plurality of viewpoints reach citizens, even as innovations might make the transmission of information more efficient, more widely shared and easier to access. A second harm can come from so-called 'fake news', which is willfully misleading information, sent for political purposes, or more generally the lack of a journalistic standard in much of the information that is available online. This scale from fake news to shoddy journalism is a matter of degree; sometimes outright lies and conspiracy theories can be found online sent for nefarious political purposes, but more often the 'fact-finding' process that characterizes journalism and news-production is just laid bare online, with all the steps being there to see for everyone. This information, as well as sometimes information discarded by professional journalists, can be both good and bad for a democracy, depending on whether it tailors to the values of pluralism, reasonableness, a competition for power and the active involvement of citizens.

A third potential problem can come from the increasing role of money in democratic politics through the increased role of the ICT environment in a parliamentary system. Technological companies are private actors whose particular corporate financial structures are often premised upon the promise to shareholders and investors of future profits. Because of this particular financing mechanism, which is much unlike the subscription-based, charitable or publicly financed models that characterize the organizations within which journalists traditionally work, profit incentives or the ability to sell advertisements can outweigh the desire to deliver facts, reasoned debate, neutrality and a plurality of viewpoints. As technological companies gain more societal influence, and thus see both their public responsibilities and private investments grow, so might the demand for realizing profits increase. This tension is inherent to the development of technology, as it thrives

through private-sector entrepreneurial competition, but ultimately impacts, and even transforms the character of, the public spaces that citizens share. So it is important to look at whether profit-incentives have usurped democratic competition between political actors, or whether democratic values such as pluralism, impartiality and reasoned debate are transforming technological corporations.

In order to address all these large and complicated issues, we adapted our research focus and methods over time. Even the conceptual clarity around these issues only crystallized as the research went on, while major events in the relation between the ICT environment and democracy continued to play out. We did from the start focus on the Dutch case, and on the 2018 Dutch municipal elections more specifically. Around these elections, we collected Twitter data from 9th of February to the 23rd of March, monitored online trackers and cookies, identified Dutch data broker and advertisement companies, kept track of the election themes, actors and events, and took stock of the major algorithms and regulatory changes that affected them. But even in this quite specific field we studied a fast-moving target, with the Cambridge Analytica Facebook scandal, the US-Russia 'fake news' scandal, the Dutch and European 'fake news' regulatory conundrums, and a whole new GDPR data law coming into effect over the chosen period. Next to that, the 'algorithms' and other design features driving the ICT environment we studied change versions rapidly, while the political environment continued to be similarly dynamic. Fortunately, this period also saw an explosion in both investigative journalism and scientific research articles that covered our field of study. When possible, we therefore restrict ourselves to our own research on the Dutch municipal elections, but will in addition often refer to research conducted by both the Dutch and the international research and journalism communities to shed light on the effect of the ICT environment on the Dutch Parliamentary system.

Parliamentary Commission Context

This specific focus on the effect of the ICT environment on the Dutch Parliamentary system was made possible with additional funding from the Dutch State Commission 'Parlementair Stelsel'. The commission has as its mission to investigate, reflect on and provide advice to the Dutch State regarding the sustainability and future viability of the tasks, functioning and position of the Dutch parliamentary system. It specifically does this in light of five developments: that Dutch citizens aspire to become more involved in politics and policy-making, that European decision-making increasingly affects both chambers of parliament, that many state capacities have been recently devolved to lower levels of government, that electoral volatility has strongly increased, and, importantly for our purposes, that the digital revolution and social media have an undeniable influence on the character of representative democracy and parliamentary institutions.

In its preliminary reports, the commission has explored several possibilities to strengthen the Dutch parliamentary system. It has done so by placing the representative organs of the state within the framework of the democratic 'rechtsstaat' (liberal democracy). A democratic 'rechtsstaat' consists of two components; a democratic and a legal component, which should be empowered in a balanced measure in order to ensure that

majoritarianism remains checked by individual and minority rights, and that the rule of law is empowered through a majoritarian democratic mandate. Central institutional components are the free and fair elections, the separation of powers, respect for basic rights and the principle of legality.

The commission has so far found that the Dutch liberal democracy generally functions well, as does the parliamentary system of two multi-party representative chambers, proportional representation and coalition-building within it. But at the same time several components require urgent renewal, also in light of an international trend of democratic decline. (Foa & Mounk, 2017; Mounk, 2018) Especially low-income and practically educated citizens are not well represented in the Dutch system, while both the constitutional resilience and legal protection of democratic institutions are insufficiently guaranteed. Technological changes also demand institutional reconfigurations. In all the this, the commission seeks to enhance the position of the voter and of citizens.

Regarding the democratic component, the commission has so far explored and provisionally recommended a binding corrective referendum, as well as the direct election of a 'formateur', who is tasked with forming a government out of a coalition of parties after elections, and who is now appointed by parliament. It also recommends more opportunities and locations for voting, as well as strengthening both the regional and the individual components of political representatives' mandates in light of their strong current national party-based affinity.

Regarding the 'rechtsstaat' component, the legal and constitutional resilience of parliamentary democracy could be enhanced by instituting a constitutional court that rules on the constitutionality of laws. It also provisionally recommends better education on democracy, sharpening the terms by which political parties may be banned, and limiting the maximum amount of political donations. Regarding the influence of the ICT environment, it provisionally recommends that political parties should be open regarding their digital campaign efforts, that party finances should be more transparent, that digital infrastructure of parties should be made more secure and that an independent public watchdog should monitor digital campaign developments. We will refine these suggestions below, especially regarding the institutional shape of the public watchdog in light of our findings, while adding some others on media plurality of our own.

The committee lastly provisionally recommends changes to ensure better interaction between the two chambers of parliament, which are the indirectly and regionally elected senate first chamber and the directly elected and proportionally represented congressional second chamber. In order to avoid either politicization or increasing irrelevance of the senate, it is suggested that it should be able to send laws back to the congress rather than have the authority to fully block them. It is also suggested that the second congressional chamber should rely more on outside information and enhance its knowledge base. We will in our final recommendation regarding a digital infrastructure for political parties connect these latter suggestions to the effect of the digital revolution on the sustainability of the Dutch parliamentary system.

Theoretical frames

Before turning to the outline of the events of the Dutch municipal elections, the regulatory environment and the political technology scandals taking place during our research period, we will provide some theoretical frameworks that have helped shape our understanding of these events, our empirical research methods, and, more indirectly, our policy recommendations. These theoretical contours relate to our two main research objects: to democracy and to the ICT environment.

Democratic Theory

Regarding democracy, within the basic framework of the democratic rechtsstaat that balances free and fair competitive majoritarian elections with basic rights and the rule of law, the work of Hannah Arendt on the relation between Truth and Politics and more recent work by Nadia Urbinati and other political scientists on the relation between populism, technocracy, media and democracy have proven particularly insightful and relevant to the topic at hand.

Arendt provides helpful categorizations of the different relations between truth and democratic politics.(Arendt, 1967) For her, scientific or rational truth is best served by an individual pursuit, while its relation to power is despotic. This means that rational, or scientific truth imposes itself, and as such is unfit to rule through the plurality of dissenting opinions that characterizes a democracy. Collective truth-finding however, to which Arendt assigns both fact-finding as well as opinion-formation, can only thrive in a democracy, as the free exchange of ideas and a multitude of witnesses is necessary to establish objective facts as well as diverging opinions. Writing about her experiences with totalitarian regimes, Arendt finds that when this collective process of the free exchange of ideas and of fact-finding is distorted, truth loses its meaning, while every non-democratic attempt to recoup it only further confuses its meaning. A desire to reinstate truth despotically can then appear, perversely backed by scientific and technological rationality, but without regaining a grip on what constitutes the free collective inquiry required to establish facts and opinions. According to Arendt, only refocusing on the human sciences and the free arts can help re-establish the delicate boundaries between truth-finding and politics, as they are uniquely equipped to make unchangeable facts bearable through imagination, and distinguishing them from those facts and opinions that require political action.

Nadia Urbinati provides a contemporary addition to what she calls the 'opinion-formation' aspect of democratic politics, which is required for, but independent of, the 'will-formation', or electoral aspect.(Urbinati, 2014) According to her, opinion-formation consists of three elements, which need to be balanced for democracy to not become 'disfigured'. The first element that is required for a democratic body of citizens to form political opinions is that the actions of those in power should be *seen* by the larger public. The second element is that every citizen should be able to *express* his or her opinion regarding those in power, and communicate it to other citizens, while the third element is that citizens should be able to *reflect* both on what they see those in power do, as well as on what opinions they hear expressed by other citizens. Now, democratic disfigurements can arise when one of these

elements becomes more prominent relative to others. When the ability to see what rulers do overwhelms the ability of citizens to reflect and express their opinions, an 'audience democracy' arises, where politicians mainly serve as entertainers for a larger public, and citizens become passive recipients of such entertainment. When the ability to express one's opinions overwhelms the ability to see what rulers do and to reflect, 'populist democracy' arises, where citizens see themselves as one body that is merely reactive in their opposition to all those 'elites' in power. When, often in an elite reaction to the other two disfigurements, the ability to reflect on power and opinion overwhelms the ability to see what is going on and express one's views, 'technocracy' arises, where expert rule limits the ability of citizens to make decisions collectively.

We have found these two theoretical frameworks by Urbinati and Arendt particularly helpful in understanding the back-and-forth between the transformation of collective opinion-formation through digitalization and the desire to control it through regulation. The attention economy that drives the advertising model behind much of the new ICT media clearly prioritizes seeing and reacting over reflecting, which can engender a technocratic counter-reaction to the extent that the new ICT environment affects the central opinion-formation institutions of a parliamentary democracy. In order to deal prudently with digital transformations, a middle ground between these has to be found, in order not to slip from Urbinati's disfigurements into an Arendtian totalitarian loss of the meaning of democratic truth.

ICT Media theory

Next to democratic theory, we have found several insights from media and cultural theory to be helpful in understanding the influence of the ICT environment on democracy in the Netherlands. The work of Geert Lovink on blogging as a technological form provides a 'determinist' technological frame, where the medium determines the message.(Lovink, 2008) For Lovink, both the technological functionalities of the ICT-environment and the views of the 'early adopters' strongly shape the kind of content that can be found on it. Technologically, internet discussions are molded by the form of the 'blog', which makes use of the 'hyperlink' structure of the net, its direct and personal nature and its very low barriers of entry to shape internet content as 'comments'. Internet content is thus due to its technological nature a reaction to other texts, and has the potential to turn news and media from 'a lecture into a conversation'.(Oudenampsen, 2018) In contrast to literary texts, where the imagination of one author can operate freely to create a whole world, or to newspaper and TV media, where editorial teams of journalists highlight particular stories to illuminate social, political and economic developments, the internet prioritizes diverse personal opinions as reactions to other viewpoints, as if whole newspapers consisted of merely the letters section, or novels of shifting personal perspectives and styles every paragraph. Though personal blogs have never become very mainstream, some of the most influential current platforms like Facebook and Twitter have copied the blog-form, where short personal posts compete for collective popularity and are opened up to commenting sections. The way in which this 'commenting' medium shapes the message on it can be seen in its 'meta'-nature, and in its tendency to react to, and question other 'mainstream'

media sources. Because the internet is a commenting system, the conversations had on it criticize traditional media sources, often in a more ironic or unserious tone. Just as print and television media reacts to and is critical of institutional power holders in a detached tone, so does online media react to and becomes critical to the third degree of other media sources. And because of its commenting nature, it is prone to continue this critical loop ad infinitum, potentially until all meaning is lost or becoming merely self-referential.

Two internet terms illuminate this mechanism. The first is the term 'circle-jerk' which puts derogatory emphasis on the tendency of discussions between like-minded persons to lead to nothing more than self-congratulation and the mutual confirmation of one another's viewpoints. The second term is that of 'Godwin's law', which states that every conversation between persons with different viewpoints on the internet would, because of its potential unlimited length, devolve into one side comparing the other side to a Nazi, or to the worst thing our collective imagination can bear. Not only can the internet's 'blogging' nature thereby devalue the value of debate itself, but it might over time also devalue the significance of collective moral pointers like Nazism, as they become terms wantonly thrown around. American white supremacist and Neo-Nazi Richard Spencer specifically credits online culture to the rise and acceptability of his ideas, stating that a large online following merely seeks to 'meme' Nazism into existence, referencing Nazi ideology in order to shock, humor and illuminate outsiders' views regarding the culture of the internet without intending to support its core beliefs(Vice News, 2016). So just as 'traditional' media seeks to break the taboos that only serve the interests of power-holders, so might the enhanced commenting nature of the internet potentially pulverize any meaning; that of the process of discussion as much as any content raised in it. Trust in shared norms and moral pointers can thereby erode.

This has several implications. Where previously comments on the actions of power-holders were mediated by editorial teams of journalists, and comments on the editorial decisions of journalists relegated to private informal conversations, both now become privileged directly in the blogging form that characterizes the internet. Media itself can thereby become politicized, as their traditional privileged commenting position is opened up, while ever more different and seemingly radical viewpoints can enter the public arena, with ever lesser seriousness attached to them.

In terms of democracy, the way the online 'blogging' medium influences its message can be both positive and negative. There is a radical democratic potential to this form because of its radically low barriers of entry, where every person becomes an equal partner in societal conversations about collective fact- and opinion-making. This optimism largely characterized the first wave of internet theories, where the liberationist potential was thought to lead to more democratic and open societies. But the flipside can be that the direct commenting nature of the internet potentially undermines the 'representative' aspect of parliamentary politics, where elected politicians, ideologues and media professionals together attempt to 'represent' a productive selection of the views of persons and groups in society, so that political conflict can become a serious clash of ideas, instead of a meaningless infinite conversational loop or a naked contest for power. This undermining

of the representative aspect of democratic tendency through the direct blogging form, and its tendency to open political contest up to a more cynical and naked competition for power has become the topic of more recent pessimistic analyses of the influence of the blogging form on media in parliamentary democracies. In our research and for our policy recommendations, we have seen both at play and advise to strike a productive balance between both.

Chapter 2. The 2018 Municipal Elections and their regulatory context

In the period that the research was conducted, the topic of the influence of the internet on democracy became of central relevance in most Western democracies. Though much information surrounding what became 'scandals' was already widely known in the research community, we could not have predicted the amount of political and media attention that took place over the past months. Three events stand out in particular.

The first was the Facebook Cambridge Analytics scandal. After an interview with former employee and 'whistle blower' Christopher Wylie of Cambridge Analytica in newspaper the Guardian on the 17th of March, the way in which big data analytics were used to profile and micro-target individuals suddenly exploded to prominence. Wylie revealed that the company he used to work for, and of which Steve Bannon, Trump's campaign manager, was one of its founders, collected user's private Facebook data and the profile data of their friends. After researcher Alexandr Kogan from Cambridge University received access to collect user's Facebook Data for research purposes and after users consented to take a personality test, Cambridge Analytica was able to get a hold of this data in order to construct profiles that were thought to be helpful for political advertising. In an unsavory possibility made possible by Facebook, not only were the data from consenting users collected, but also those of all their non-consenting friends. And in an especially shady move by Cambridge Analytica, information collected for academic research purposes was commercialized and used for personally targeted political campaigning.

This monetization of personal data and it being used in political elections was premised on the presupposition that a link existed between behavioral personality profiles and political preference, as US campaigners have long believed to be the case. By linking, through small tests, personality profiles to Facebook profile data points, Cambridge Analytica, and with it the Trump campaign, thought it possible to predict voting behavior based on Facebook data. With over 87 million Facebook profiles collected they also hoped to be able to identify and target groups essential to their electoral strategy both on- and off-line, and target them messages tailored to their personal data.

The whistleblower interview and the following media outrage led to parliamentary hearings in both the US Congress and the European Parliament of Facebook's president Mark Zuckerberg. He was able to convince politicians that the leaks were a one-time affair, and that he would increase investments into security and privacy[1]. He also promised he would open up political 'dark' ads used on the platform for everyone to see, and to refocus on the sharing of personal information through Facebook blogging, away from the political and

---

[1] This proved in vain; on the 28th of September Facebook announced another 50 million profiles were leaked due to a hack. (https://bgr.com/2018/09/28/facebook-data-breach-2018-yep-another-one/)

news content the platform had become known for. Though Zuckerberg was thought to get away quite easily with massive privacy violations in the parliamentary hearings, Facebook stock took a 19.6%, 123.4 billion dollar plunge because it chose to forgo profits in favor of massive increases in privacy and security investments, and enhancing the quality, instead of quantity spent on the platform (Dillet, 2018).

The impact of the Cambridge Analytica and Facebook leak on the Netherlands was not as big. Of the 87 million leaked profiles, a maximum of 90 000 Dutch user profiles were possibly improperly used by Cambridge Analytica. But political micro-targeting techniques similar to those employed by Cambridge Analytica have been used by almost all political parties in the Netherlands during the municipal elections, and are on offer from a host of data brokers, which we identify in chapter 3. Facebook in turn is offering the same micro-targeting capabilities in the Netherlands as anywhere else, and again most political parties have admitted to making use of Facebook's to target potential voters.

The second event that brought the topic of the influence of the internet on parliamentary democracy to prominence during our research period concerned the perceived threat of Russian disinformation. As many states do, through information campaigns the Russian state seeks to put the actions of its own actions in a positive light. It has done this where it has been accused of geopolitical misconduct, such as in its interference in Ukraine or Syria, its occupation of the Krim, or surrounding the crash of Dutch flight MH17 and the poisoning of Russian nationals in the UK. But Russian foreign policy 'propaganda' has extended beyond this in that it sought to rather than support one side than another, simply polarize political opinions in western democracy in order to confuse and incite publics, and enhance the reach of 'disinformation'. On the website 'EU vs. Disinfo', which monitors misleading reports pushed by Kremlin sources, false stories regarding migrants have been reported according to which the EU is both building concentration camps for migrants, as well as not prosecuting any crimes committed by migrants. In the US, Russian sources have created Facebook groups and organized rallies for opposing sides, pitting white nationalists against anti-fascists in on- and off-line face-offs that would not have taken place otherwise. Russian policy seems intent on bringing Western trust in political and economic institutions of liberal democracy down to the levels of Russia itself.

In the Netherlands, our research has not found a lot of evidence of Russian online interference, measured in terms of the distant proxy of bot activity, surrounding the elections. Investigation from journalists at NRC found that over 900 Russian accounts were active in the Netherlands in 2017, but that these were not active during the Dutch 2017 elections. Their messages focused instead on anti-islam propaganda, and their activity peaked surrounding the US elections and major terrorist attacks. Far from merely pushing their own stories, they amplified the views of right-wing nationalist bloggers and twitter users, in order to spread the reach and effect of disinformation, lower trust in shared institutions and policies and increase polarization. But a particular interest in influencing Dutch elections specifically is hard to delineate.

The third event that arose during our research concerned the way in which disinformation can breed more disinformation, witnessed in the actions of the aforementioned European Commission's EU vs. Disinfo group took against Dutch media. Housed under and funded by the European External Action Service, the EU vs. Disinfo service is tasked with identifying and highlighting 'one-sided information, incorrect facts, or disinformation coming from or being in line with messages put forth by the Kremlin'. But in a concerning instance of over-reach, articles, blog-post and radio-items of the Dutch public radio news channel, of regional paper 'de Gelderlander', and of right-wing blog 'TPO' and shock-blog 'Geenstijl' were in January identified by the EU vs. Disinfo group as 'disinformation'. In each and every one of these cases, the EU group had to retract its conclusions. Lacking any native Dutch speakers, in some articles the group had mistaken quotes that reporters got at political rallies with stated journalistic facts. Though the lower chamber of Dutch parliament with a majority rejected continued funding for the EU vs. Disinfo group, the Dutch minister of internal affairs Kajsa Ollongren continued to support them, after having previously warned about the threat of Russian 'fake news' in the Netherlands. And as our analysis of Twitter data shows, in a strange turn of events this central 'fake news' theme became one of the main themes of the Dutch municipal elections.

The 2018 Dutch municipal elections

The Dutch 2018 municipal elections took place on the 21st of March in 335 municipalities. The turnout was with 55% much lower than the previous national elections (82%), but about average compared to other municipal elections. The campaign preceding it was relatively unremarkable, while the outcomes were more remarkable in that they revealed strong wins for new and local parties. We chose to focus on the city of the Hague in addition to the national frame. In the Hague over 15 parties participated, with a strong and ideologically eclectic mix of national and local parties. The city is, as all three major Dutch cities, characterized by hyper-diversity, with around 50% of residents being first- or second-generation immigrant. The city is known both for being one of the Netherlands' central terrorist breeding grounds, as well as for being the location of many international organizations as well as the institutions of the national government. Prior to the 2018 elections, the Hague was governed by a diverse array of parties; the social-democrat PvdA, the liberal-democrat D66, the liberal VVD, the Christian-democrat CDA and the local party HSP, which, together with populist-right PVV represented the main winners of the 2014 elections.

Election highlights, election outcomes

The 2018 campaign preceding the election was characterized by several, largely scattered, debates and events. The first debate in Amsterdam between leaders of national parties took place on February 9th, and had as its main themes besides housing, public transport and Schiphol airport also several escalations related to racism and identity. After a local politician from newcomer party FvD maintained in an interview that there was a relation between race and IQ, several other parties took turns attacking the FvD party leader on its purported racism. On the 13th of February the minister of Foreign Affairs of liberal party VVD

was caught in a lie and resigned, and during the election minister Ollongren of internal affairs became embroiled in a fierce debate surrounding 'fake news' mentioned above. Two days before the elections, final debates were held at both the national and local levels. For the Hague, the themes revolved around housing and camera surveillance. At the final national debate, affordable housing was also a central theme, but with more explicit ideological links to the role of migration in housing shortage versus the role of shareholder tax deductions in shrinking housing expenditures. Language requirements for migrants and other 'integration' related themes also played a central role.

The election outcomes at the national level showed strong wins for local parties, while the Liberal VVD dethroned the Christian-Democrats as the biggest local party. The national ranking of parties tailors well with the previous Dutch national elections, where a coalition of D66, CDA, VVD and Christian Union govern together. The strong showing of local parties is usually interpreted as a sign that local democracy is working in municipal elections, as different local issues are represented differently in different places. Due to the large overlap with national outcomes, there doesn't seem to be a large discrepancy between the positions that voters have expressed to local versus national political parties. But having a link to national party

Party	Votes		Seats	
	#	%	#	+/−
Local Parties	1,945,904	28.65	2,566	▼ 36
People's Party for Freedom and Democracy	916,930	13.50	1,117	▲ 19
Christian Democratic Appeal	910,441	13.41	1,277	▼ 221
Democrats 66	623,671	9.18	594	▼ 231
GroenLinks	602,157	8.87	520	▲ 166
Labour Party	510,527	7.52	544	▼ 255
Socialist Party	301,600	4.44	282	▼ 158
Christian Union	260,007	3.83	300	▼ 59
Other Parties and combined lists	720,467	10.52	686	▲ 91
	6,849,788 [6]	100.00	7886	▼ 684

might actually work against a local party's ability to accurately represent local issues, or local chapters feel it against their interest to own up to the national policies of the mother party. This round, several more local chapters chose to go their own way, and separate from the local party. The desire of strong national parties like PVV or FvD to expand locally was also not very successful; often local competitors with similar programs won out, or local affiliated chapters chose to sever links with the national party altogether.

In the Hague, the local party of Richard de Mos outperformed all polls and became the biggest party. De Mos, a former PVV MP, conjured together an eclectic mix of ideologies and personalities. De Mos previously profiled himself as a classic 'right-wing populist' politician, who 'fights back' against the changes to Dutch society that multiculturalism and hypermigration bring, favoring secularist re-education for misbehaving youth, more money for elderly care and strong protection of traditional Dutch values. But midway through his previous term he recruited D66 politician Rachid Guernaoui, whose party often positions itself as the liberal cosmopolitan opposite of the populist right. De Mos meanwhile began to take on almost every issue, positioning himself has a pragmatic local 'ombudsman'. His party sided with the Green party on most green issues, and tried to appeal explicitly to the

diverse array of minorities in the Hague. This is why, during the campaign, we to our surprise found smaller parties traditionally representing both sides of Dutch identity and migration politics such as PVV (far right populist) NIDA (progressive islamist), DENK, Islam Democrats and Party of Unity (Pro-Islam, other minorities) all fretting over de Mos taking their votes. With his pragmatic, local and populist approach, de Mos was able to bring together working class voters from migrant as well as Dutch backgrounds and become the biggest party in the Hague.

Party	Votes		Seats	
	#	%	#	+/–
Group de Mos / Heart for The Hague	32,632	16.80	8	▲ 5
People's Party for Freedom and Democracy	26,990	13.90	7	▲ 3
Democrats 66	23,973	12.34	6	▼ 2
GroenLinks	20,684	10.65	5	▲ 3
Christian Democratic Appeal	12,618	6.50	3	— 0
Labour Party	12,567	6.47	3	▼ 3
The Hague City Party	10,897	5.61	3	▼ 2
Party for the Animals	10,201	5.25	2	▲ 1
Party for Freedom	9,036	4.65	2	▼ 5
Islam Democrats	6,943	3.57	1	▼ 1
Christian Union – Reformed Political Party	6,129	3.16	1	— 0
Socialist Party	5,176	2.66	1	▼ 1
NIDA	4,481	2.31	1	New
50PLUS	4,444	2.29	1	New
Party of Unity	4,026	2.07	1	— 0

Regulatory environment in the Netherlands

During the research period, the Dutch law on protection of personal data was replaced by the European General Data Protection Regulation (GDPR) on May 25th. Though in both regulations, the protection of personal data was already quite extensive compared to the US context, the GDPR has further extended the protection of personal data. GDPR imposes obligations on organizations that process personal data, and grants rights to people whose data is being processed, for example to see how an organisation is using one's data. (Zuiderveen Borgesius et al., 2018) In the Netherlands, the independent supervisory authority 'Autoriteit Persoonsgegevens' is tasked with overseeing compliance with the rules. In addition to general personal data processing, there are additional restrictions on the use of special categories of data, such as political data. In article 9 of GDPR, the processing of personal data revealing political opinions is explicitly prohibited, but there are quite some exceptions. The most important exception is that personal data may be processed when a subject has given explicit consent. Others are that an organization (like a political party) may process data related to political opinions when that person is a member, or that data

may be processed when a person has already manifestly made his or her political opinion public.

Data processers must also be transparent about which data they collect and the purpose for which they collect data, which they must disclose in for example a privacy statement. GDPR also requires companies to disclose the use of trackers and cookies.

As with much new regulation, the devil will be in the details, and repeated data request inquiries of persons as well as litigation will have to refine the terms on which organizations are and are not allowed to process personal political data. As of now, only very few organizations yet comply with these new GDPR regulations, but a period is granted by authorities for companies to comply.

In addition to personal data protection laws in the Netherlands the advertising industry has established a sectoral code for social media advertising, the 'Reclame code commissie sociale media'. That code also specifies that advertisers should provide transparency, and also includes a prohibition on manipulation. Political advertising is however not yet explicitly addressed, though it can be thought to be included.

It is lastly, and fortunately, very hard to restrict 'political speech' in all western liberal democracies, protecting political actors from potential infringements to privacy violations. In the Netherlands, article 10 of the Constitution protects individual privacy, but case law at even the European level has found that this provides little scope for restrictions in political speech. Especially political parties' speech cannot be restricted, due to their essential role in fostering open debate in democratic elections. But whether data brokers and other companies that have now become necessary intermediaries for the expressions of the speech of political parties also fall under this protection is in Europe and the Netherlands still largely an open question. Less so in the United States, where companies have controversially been protected in their financial contributions to politicians under the header of free political speech.

Chapter 3. Political Micro-targeting in the Dutch Parliamentary System

In order to answer our Sub-RQ2, "What particular tracking mechanisms (companies and websites) create an online advertising environment conducive to filter bubbles and echo chambers?" we looked at the role of big data research and individual tracking mechanisms in the 2018 Dutch municipal election. Together with the infrastructure provided by technological corporation such as Google and Facebook, these together make political micro-targeting possible. In this chapter, we will set out what political micro-targeting precisely is, discuss its essential components, and investigate their role in the 2018 elections. We also provide insights into the underlying mechanisms driving micro-targeting, and the role played by regulatory and economic incentives. This helps us to give a final set of policy recommendations on how to regulate the ICT environment so that it enhances Dutch democracy.

So what is political micro-targeting, and what are its essential components? Political micro-targeting allows political actors to target individual citizens directly in their online browsing, so that every user gets a different, specifically tailored advertisement to see. As it stands, there are three necessary components to make political micro-targeting function, which this chapter will discuss in turn.

1. An online infrastructure that presents advertisements and creates a revenue model for doing so,

2. Individual data brokers, trackers and companies that compile personal or categorical profiles

3. Big data models linking voter or other political data to all other data points collected under 2.

The first and second functions are currently largely taken up by two companies: Facebook and Google. These two showed up most in our inventory of trackers on Dutch political sites. But a host of smaller companies are also making use of the advertising infrastructure provided by Google and Facebook to compile profiles through trackers, and to create big data models for the expression of political opinions of voters.

Together, these three necessary components allow political parties to micro-target to an ever greater and more precise extent. Cambridge Analytica, but also to lesser degree more 'progressive' companies such as Blue State Digital use the combination of the online advertising infrastructure, personal profiles and political models to target voters. This raises problems for the health of parliamentary democracy to the extent that this 'datafication' hampers the ability of citizens to engage in 'self-legislation', The ideal of democracy is that citizens are able to make the laws for themselves, or 'self-legislate'. This ideal of self-

legislation goes deeper than democracy: it also lies at the basis of the ideals of freedom and autonomy, where one should be able to set the terms upon which one seeks to live for oneself and engage with others.

The problem with political micro-targeting might thus be that persons become entangled in an ever-expanding web of data representations held by private or governmental actors, that are then used to influence personal behavior. Privacy is violated, voters manipulated, and only those groups are targeted that are key to swinging elections or particular business interests. Power might also be transferred to commercial intermediaries, while politicians and data brokers can further shield themselves from declaring their intentions. If done well, the enhanced technological means behind political micro-targeting might however also lead to better self-legislation, diversifying the number of political options available to citizens, reaching new socially excluded groups and making political communication more effective. (Zuiderveen Borgesius et al., 2018)

The primary focus of regulation should thus be the extent to which political micro-targeting is hidden or opaque; not just because researchers and policy-makers would like to better inform themselves, but because knowledge about the factors influencing one's behavior is essential to the extent to which citizens are able to be free and democratic self-legislating citizens that can make informed decisions.

The online advertising infrastructure

The biggest problem facing our research, as well as that of any other, has been that big players like Facebook or Google, but also smaller data brokers, have not been willing to show the way in which they make use of the components required for political micro-targeting. Because of that, it is impossible to investigate the extent to which they violate privacy directives, and seen from the democratic ideal of 'self-legislation', this is a problem in itself. This became especially clear in the public and political debates surrounding the Cambridge Analytica scandal, where neither of the three essential components required for political micro-targeting received much attention. Facebook has since commenced a procedure in which researchers, after careful review, can access data for research, though at the time our research commenced no such cooperation was forthcoming.

Even though the Cambridge Analytica scandal and Facebook's political hearings brought political micro-targeting to the public's attention, the essential mechanisms driving the infrastructure behind it remained underexamined in de media-storm surrounding the 'leaks' from Facebook. What remained unclear is that same principles that outraged the public regarding the leaks drive all of Facebook's business model, and any big data analytics company.

Facebook's own internal targeted 'advertising platform' is able to generate revenue because it allows sellers to promote their products to targeted but anonymized groups. To advertisers, Facebook is selling over 29 000 unique categories it collects on its users, and it has admitted to have bought and used data from many different sources other than the ones provided by users on its platform. One such category can for example be 'white boat-

loving classical music listening former soccer playing women', with the categories becoming so fine-grained that the targeted group can be made as small 10 people. This categorical targeting can violate personal data directives when one is willing to generate 9 fake profiles, to sort out the one.

Behind these targeted advertising possibilities are models based on partial but often gigantic data-sets, collected by online trackers and cookies, as well as public data-sets. Many different companies collect this data, and Facebook has developed its own hidden 'Pixel' trackers that companies are encouraged to install on their own websites in order to collect data on their visitors and track them online.

The models that Facebook and other data companies have developed are iteratively refined by feeding more new data into it. Newer iterations of these models are now also able to 'self-learn', or to automate the process of representing data for specific purposes, and seeking out more data to refine its focus on that purpose. By 'anonymizing' the data fed into it, for example by replacing names with numbers, personal privacy violations are pre-emptively believed to disappear, even as one still receives ads that were targeted to one's identity as it differentiates you from billions of other people. In other words, both Facebook and other data companies have thus constructed 'profiles', or representations of almost every single person (even those not on, or consented to their services) and it is the models that refine and collect these representations that are what makes the company valuable. So even when Facebook or other companies vow to not 'store' or 'own' the information one puts on it, that information is still likely to have been used to refine the identifiers in the models that are used to target you. Moreover, once the model has refined its classification of any one person, it can discard the personal data once used, as it has 'learned' them for purposes of targeted advertising.

The whole concept of individual privacy thus becomes warped through these developments, making it ever harder to identify what data belongs to whom. So it also goes for explicit consent; as consent given once, having served to refine the categorization of the model, is almost impossible to meaningfully turn back. As data brokers and Facebook itself furthermore operates behind other companies, not giving consent to 9 out of 10 food company websites but giving it to 1 company a person trusts suffices to add one's food preferences to the model constructed and sold back to any different company or organization one has consented to. In this way, it becomes increasingly difficult for person to engage in meaningful 'self-legislation', when the information asymmetry between data collectors and data subjects is this big. And for data companies, like Facebook, it becomes increasingly easier to deny knowledge of the privacy or consent points in their own data models, as Zuckerberg repeatedly did in the hearings.

Dutch trackers active on political sites

As it was impossible to collect data on the models used by Facebook and other data brokers, which are the key pieces of intellectual property that give value to the companies, we instead inventoried the trackers present on relevant Dutch websites. These relevant Dutch websites consisted of all major news sites, the sites of all political parties, both at the

Hague and national level, and a long list of clickbait, 'fake news' and personal blogging sites, retrieved from an helpful online source (https://sites.google.com/site/dehoaxwijzer/valse-nieuwssites) and from personal browsing. We have inventoried the trackers for all sites, as well as for the sites of Dutch political parties only.

For all trackers, we have created a visual representation that reveals the links between them, eradicates duplicates, and shows through the trackers' respective sizes how many other trackers are present when the one is present.
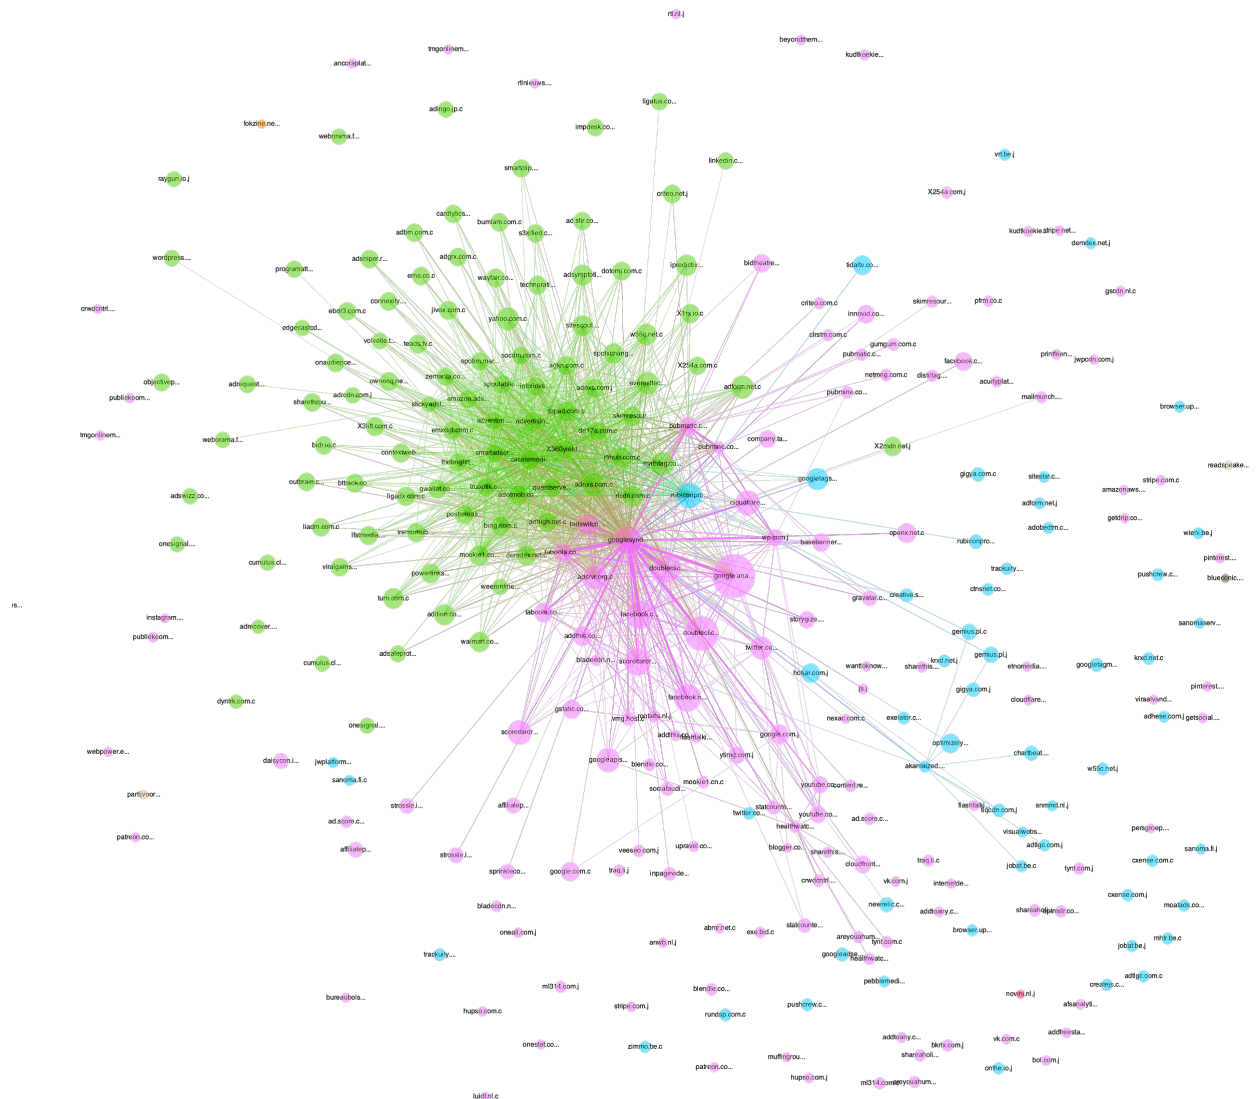
The assumption underlying this analysis is that methods propagate among allied sites. If I uses some trackers and you either hire me to set up a site for you or ask me what to recommend, I'll give you mostly the same trackers that I use.

22

If this is right, then we can see three main communities of third-party kit, pictured in green, purple, and blue. There are actually 9 modularity classes in this dataset, but 6 of them are very tiny. The purple kit is dominated by Google, along with Facebook and some others. The green kit are primarily trackers and targeters. We looked up a few of the bigger nodes and found companies such as:

* getintent.com: "The only AI-driven, next generation bidding engine. Unmatched flexibility: unleash your data. Your algorithms. Your success."

* Mediamath.com: "MediaMath collects and processes personal data about you via the Platform, when you visit the Website, and when we engage in marketing activities based on contact information you provide to MediaMath through other means. [...] MediaMath is a global provider of digital advertising media and data management technology. We provide a technology platform and services which help advertisers do a more effective job of reaching their targeting audience.

* zemanta.com: "The Native DSP: Stunningly Intelligent Native Ad Technology"

* truoptik.com: "The leader in OTT TV measurement, data management, targeting, and attribution. 70 Million US households identified, 10,000 audience segments, deterministic behavioral data across over 300,000,000 homes globally"

The blue kit is like the green: it seems to be mostly micro-targeting sites. We don't know enough about this area to say whether any of these third-party domains are malicious. But we have plenty of evidence now that the websites in question are engaging in some sort of micro targeting.

Regarding specifically the trackers used on the sites of political parties, the number of trackers was much lower, as the below image reveals. Only one party, the Haagse stadspartij, used a tracker known for its malicious software.
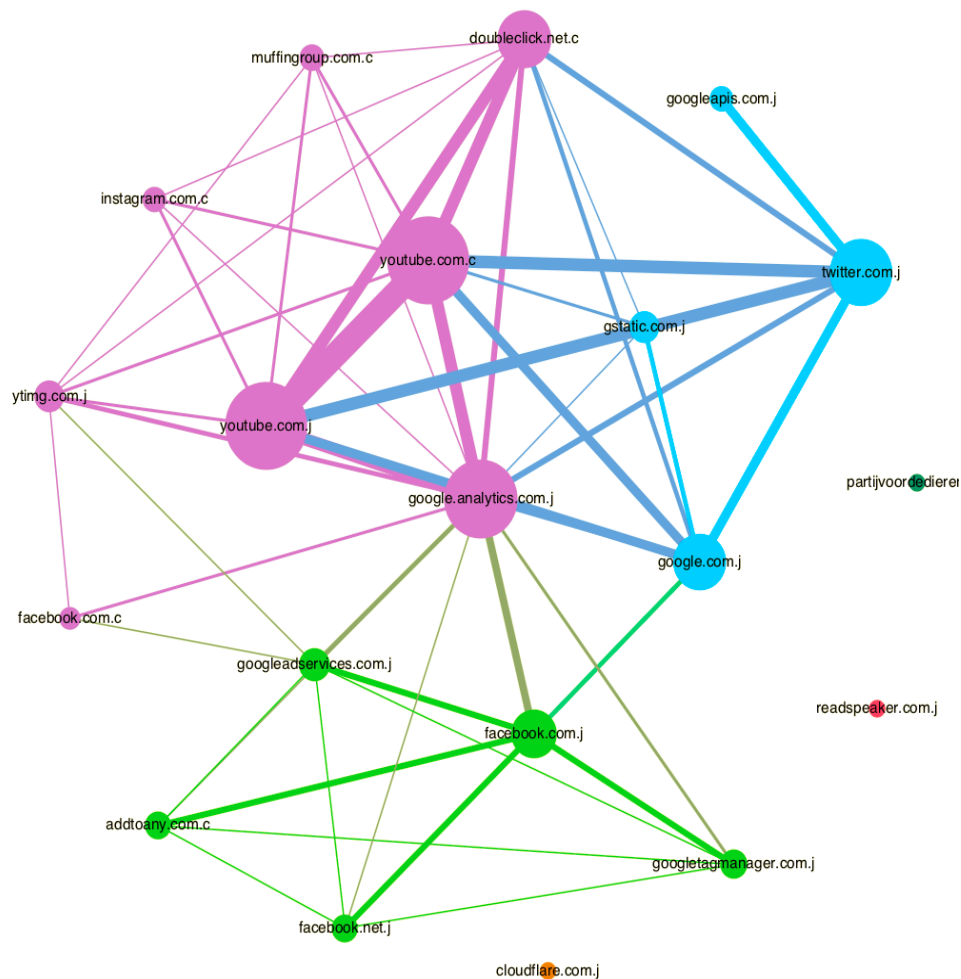
It thus seems major parties do not themselves engage in large-scale micro-targeting. This however says nothing about them hiring the services of other data broker companies, or making use of Facebook's advertising platform. But before turning to those, we have early in our research also duplicated the work done by such companies, in order to show that it is possible with limited means to construct big data models required for political micro-targeting in the Hague.

Big Data analysis of demographics and voting behavior in The Hague

In order to show that the third component of political micro-targeting can work in the Hague we have analyzed how demographic characteristics of the citizens can "predict" voting behavior. We have tried to replicate the process by which one constructs a model on the basis of data gathered by individual trackers and cookies and publicly available datasets, which can in turn be used to 'canvass' door to door or select one's target audience in Facebook's micro-targeting platform. By demographic characteristics, we mean age,

ethnicity, income, employment, etc. In case of voting behavior, it is a number of votes given for particular party, for example, PVV, PvdA, D66, and so on. The analysis was conducted for The Hague on the district level. That is, 44 districts were examined. As voting datasets municipal elections of 2014 and 2018 were chosen. As main instruments 3 machine learning algorithms were used - NMF, multiple linear regression (MLR), and support vector regression (SVR). All manipulations were made in Python programming language under support of Scikit-learn library.

In short, the job was done in the following manner. First, the dataset for the 2014 year was analyzed. 33 demographics characteristics were carefully chosen to "cluster" citizens using NMF. As a result, 7 clusters-profiles were identified (Figure 1). The interpretation here is as follows. Let us consider column number 3 ("profile" number 3). Here we can observe grouping of old Dutch citizens with low-income received from pensions. Secondly, these profiles were matched to 44 districts (Figure 2). This figure can be read in the same manner. That is, previously mentioned profile number 3 can be mostly found in "Zuiderpark" district. These results were roughly validated by the citizens of The Hague.

When this model would for example be used for door-to-door campaigning, people responsible for interviews are provided with a "map" that shows where they can find people with particular demographic characteristics. In addition, results were plotted on the map (see Jupyter notebook).

The same procedure was then applied to voting data (Figure 3 and 4), 7 voting profiles were identified and aligned to the districts. The voting profiles here represent unique combinations of votes for various parties, so they do not yet tailor specifically to parties but rather to 'electoral clusters'.

Finally, we combined obtained results. Using voting profiles as predictive variables and a set of corresponding values for demographics profiles as predictors. That is, MLR and SVR were applied to get "predictions". As a metrics to evaluate results R-Squared was chosen. In case of MLR, the results can hardly be named as "satisfactory", since only 3 out of 7 voting profiles can be predicted with "accuracy" (R-Squared value) from 0.53 up to 0.75. But SVR was much more satisfactory; provides R-squared values from 0.75 to 0.86 for 6 out of 7 profiles.

Important to mention some limitations of the conducted analysis. First, assumptions of the methods used should be studied more carefully. Although, using NMF as a clustering algorithm allows to partly avoid multicollinearity (import for the adequate usage of MLR) or SVR, in its turn, is less strict then MLR and has non-linear kernels. Secondly, to see more clearly the connection between demographics and voting behavior, recent datasets in demographics should be studied. Currently some data, income values, for example, available only for 2014.

To conclude, it seems that microtargeting might work in The Hague. Using latest knowledge in machine learning and publicly available datasets it is possible to get satisfactory estimates. Therefore, decision-makers should be aware of such technologies.
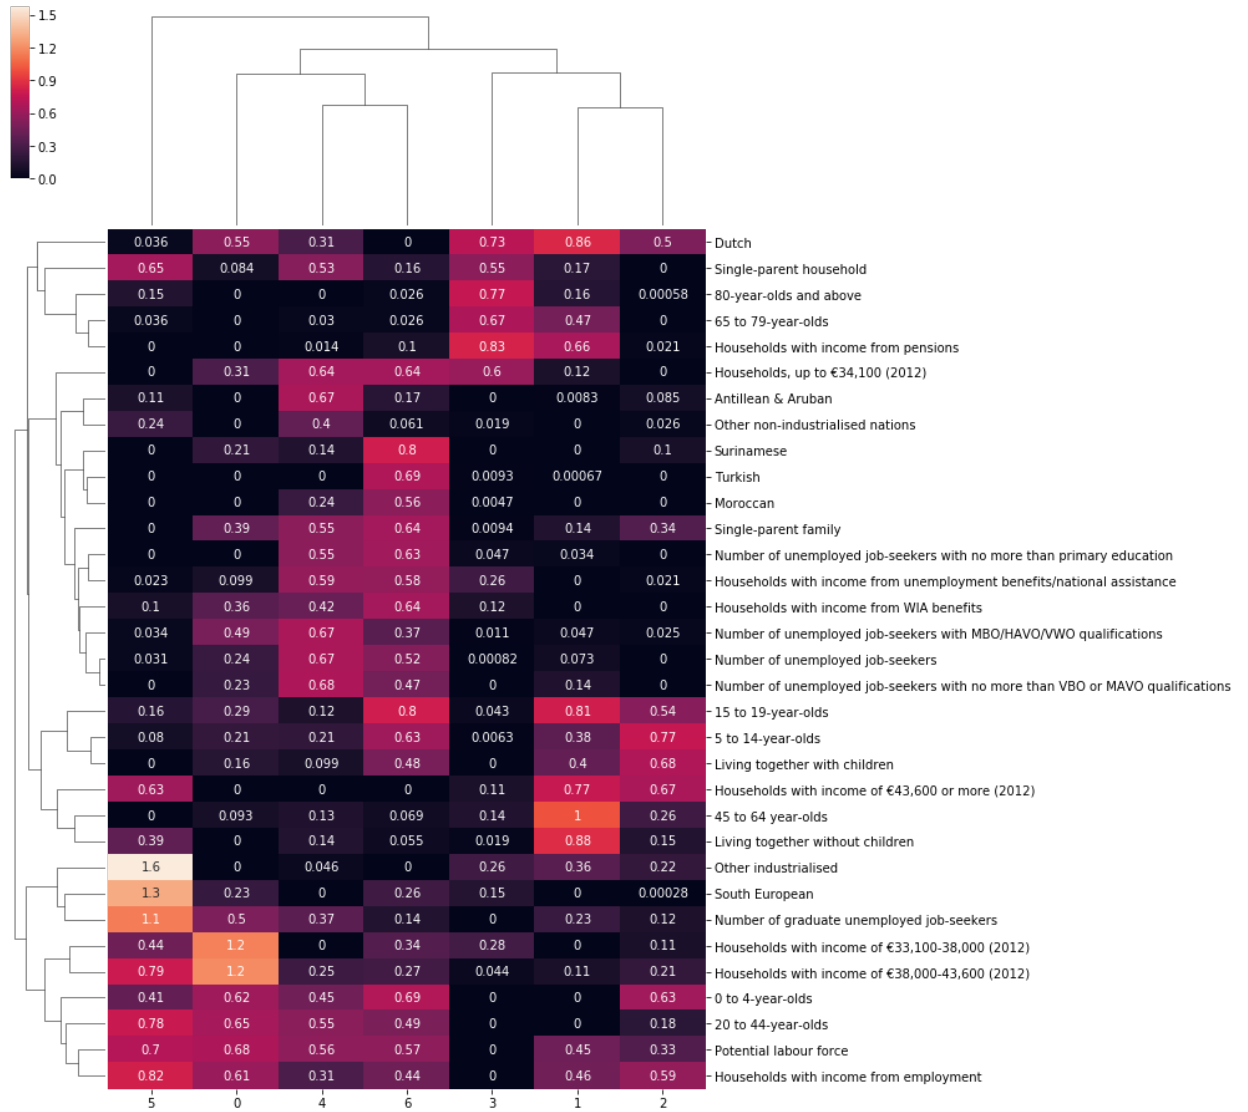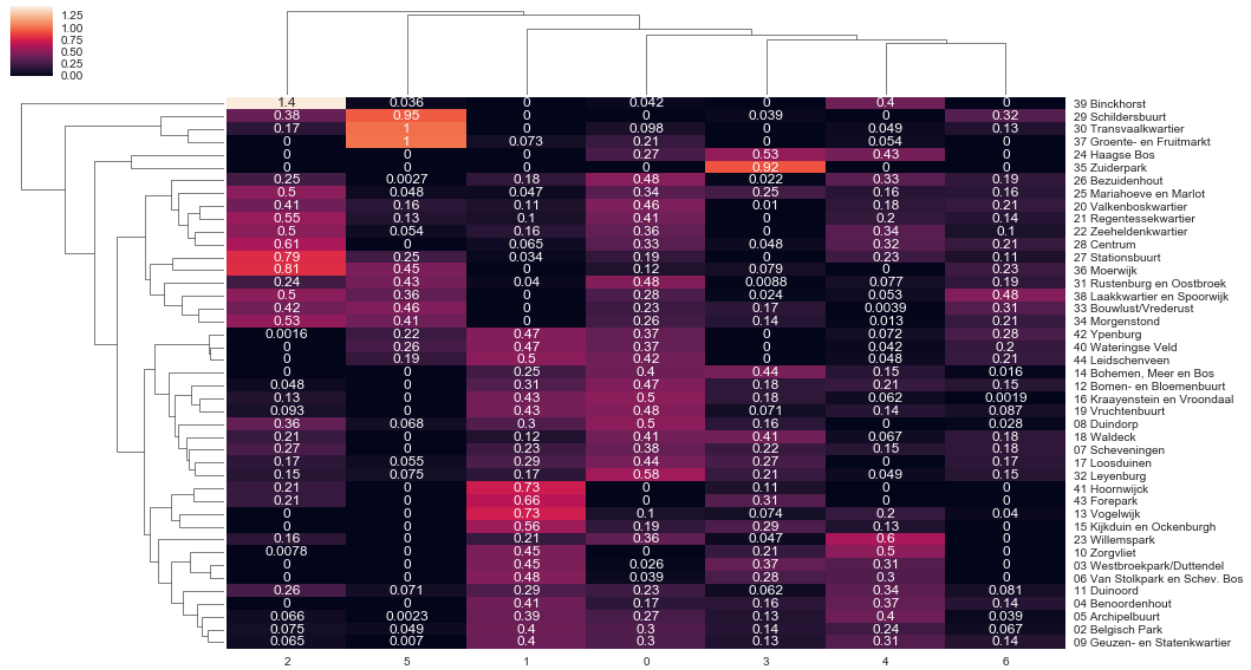
Figure 1. "Clustering" voting preferences into profiles

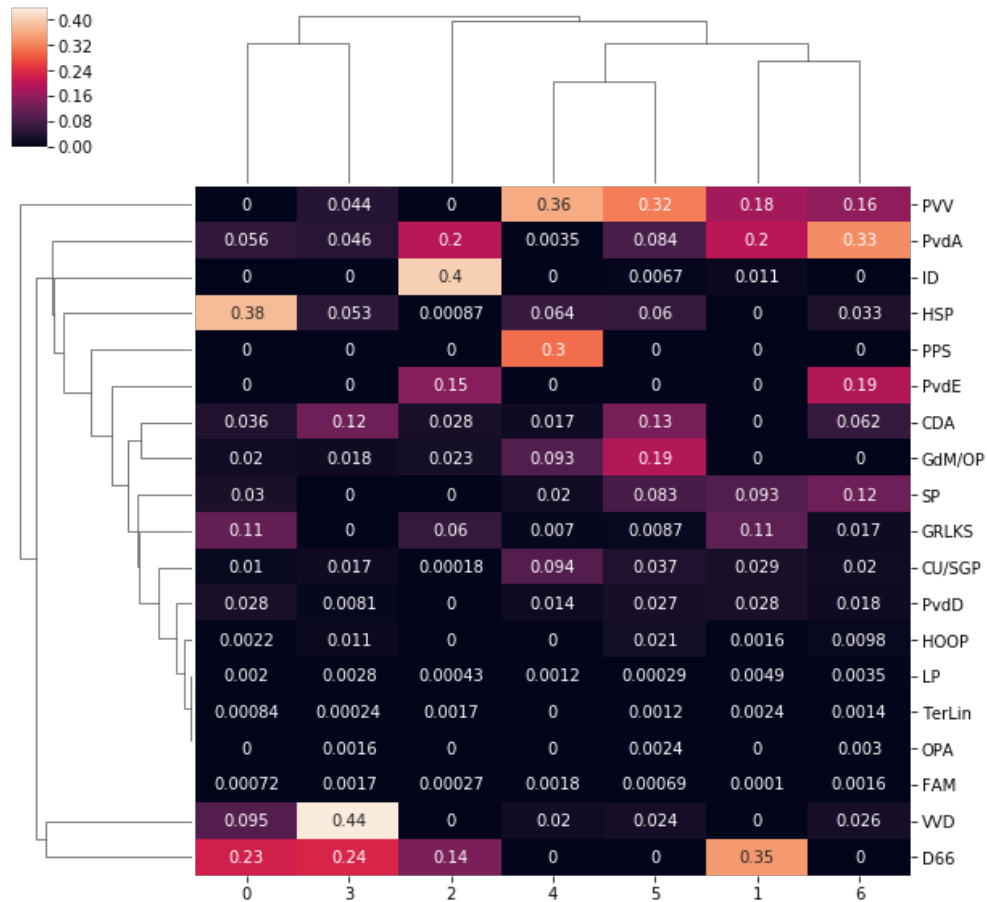Figure 2. Distribution of demographics profiles across the districts

Figure 3. "Clustering" voting preferences into profiles

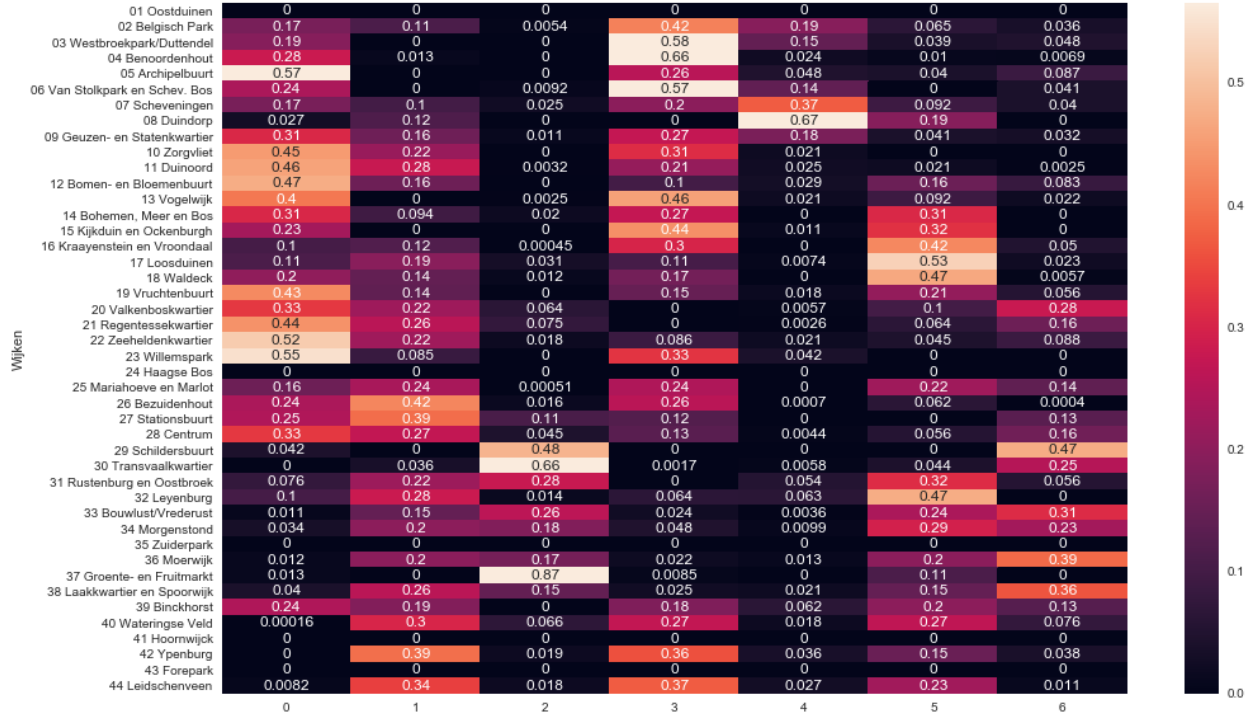| Wijken	0	1	2	3	4	5	6
01 Oostduinen	0	0	0	0	0	0	0
02 Belgisch Park	0.17	0.11	0.0054	0.42	0.19	0.065	0.036
03 Westbroekpark/Duttendel	0.19	0	0	0.58	0.15	0.039	0.048
04 Benoordenhout	0.28	0.013	0	0.66	0.024	0.01	0.0069
05 Archipelbuurt	0.57	0	0	0.26	0.048	0.04	0.087
06 Van Stolkpark en Schev. Bos	0.24	0	0.0092	0.57	0.14	0	0.041
07 Scheveningen	0.17	0.1	0.025	0.2	0.37	0.092	0.04
08 Duindorp	0.027	0.12	0	0	0.67	0.19	0
09 Geuzen- en Statenkwartier	0.31	0.16	0.011	0.27	0.18	0.041	0.032
10 Zorgvliet	0.45	0.22	0	0.31	0.021	0	0
11 Duinoord	0.46	0.28	0.0032	0.21	0.025	0.021	0.0025
12 Bomen- en Bloembuurt	0.47	0.16	0	0.1	0.029	0.16	0.083
13 Vogelwijk	0.4	0	0.0025	0.46	0.021	0.092	0.022
14 Bohemen, Meer en Bos	0.31	0.094	0.02	0.27	0	0.31	0
15 Kijkduin en Ockenburgh	0.23	0	0.00045	0.44	0.011	0.32	0
16 Kraayenstein en Vroondaal	0.1	0.12	0.00045	0.3	0	0.42	0.05
17 Loosduinen	0.11	0.19	0.031	0.11	0.0074	0.53	0.023
18 Waldeck	0.2	0.14	0.012	0.17	0	0.47	0.0057
19 Vruchtenbuurt	0.43	0.14	0	0.15	0.018	0.21	0.056
20 Valkenboskwartier	0.33	0.22	0.064	0	0.0057	0.1	0.28
21 Regentessekwartier	0.44	0.26	0.075	0	0.0026	0.064	0.16
22 Zeeheldenkwartier	0.52	0.22	0.018	0.086	0.021	0.045	0.088
23 Willemspark	0.55	0.085	0	0.33	0.042	0	0
24 Haagse Bos	0	0	0	0	0	0	0
25 Mariahoeve en Marlot	0.16	0.24	0.00051	0.24	0	0.22	0.14
26 Bezuidenhout	0.24	0.42	0.016	0.26	0.0007	0.062	0.0004
27 Stationsbuurt	0.25	0.39	0.11	0.12	0	0.056	0.13
28 Centrum	0.33	0.27	0.045	0.13	0.0044	0.056	0.16
29 Schildersbuurt	0.042	0	0.48	0	0	0	0.47
30 Transvaalkwartier	0	0.036	0.66	0.0017	0.0058	0.044	0.25
31 Rustenburg en Oostbroek	0.076	0.22	0.28	0	0.054	0.32	0.056
32 Leyenburg	0.1	0.28	0.014	0.064	0.063	0.47	0
33 Bouwlust/Vrederust	0.011	0.15	0.26	0.024	0.0036	0.24	0.31
34 Morgenstond	0.034	0.2	0.18	0.048	0.0099	0.29	0.23
35 Zuiderpark	0	0	0	0	0	0	0
36 Moerwijk	0.012	0.2	0.17	0.022	0.013	0.2	0.39
37 Groente- en Fruitmarkt	0.013	0	0.87	0.0085	0	0.11	0
38 Laakkwartier en Spoorwijk	0.04	0.26	0.15	0.025	0.021	0.15	0.36
39 Binckhorst	0.24	0.19	0	0.18	0.062	0.2	0.13
40 Wateringse Veld	0.00016	0.3	0.066	0.27	0.018	0.27	0.076
41 Hoornwijck	0	0	0	0	0	0	0
42 Ypenburg	0	0.39	0.019	0.36	0.036	0.15	0.038
43 Forepark	0	0	0	0	0	0	0
44 Leidschenveen	0.0082	0.34	0.018	0.37	0.027	0.23	0.011

Figure 4. Distribution of voting profiles across the districts

What we thus for example read in the results, in figure 3, is that D66 voters overlap to a great degree with HSP or VVD or PvdA and PVV voters; so that D66 politicians seeking to win votes should target voters from these three other parties. Where these 0, 3 and 1 cluster voters live can then be seen in figure 4, and represented on a map. The identity of the voters that make up these clusters can then be seen in figure 1, so that D66 would know it should target elderly Dutch voters, Dutch voters with high incomes, as well as Dutch voters with low incomes. Targeting young migrants in the Hague is relatively fruitless, even though the ideological inclinations of D66 might push them to do so. On Facebook, it is easy to tailor political advertisements to these categories.

We also see that the model was not able to meaningfully assign Groep de Mos to specific categories. This shows that our limited political micro-targeting model was not able to predict the rise of de Mos, as most polling also failed to do. For the health of democracy overall, and the ability of Dutch parties and citizens to 'self-legislate' outside of the apparent constraints of political micro-targeting this can be seen as positive news. But due to the relatively limited scope of our model, which has not relied on personal data profiles built by tracking companies, and due to reports of the worrying secrecy surrounding the funding given to de Mos for campaigning, these optimistic speculations cannot yet serve as conclusive.

Micro-targeting capabilities in the Netherlands and their effect

We thus find that political micro-targeting is both possible and probable in the Dutch parliamentary system. These findings are confirmed by newspaper reports and interviews on political micro-targeting in the Netherlands. Especially the firm 'de Politieke Academie' has stated that it makes use of all three components of micro-targeting and serves several political parties in the political campaigning. By combining over a hundred different categories with voter data, the firm has stated it is able to bring voting data, which restricts data to clusters of 1200 people, to smaller geographical clusters of 35 people. Between 70 and 80 political parties have hired the services of this political micro-targeting firm alone this municipal election. (Bouwman, 2018) All parties also admit to make use of Facebook's targeted advertising possibilities.

We can thus see where the regulation here might go awry of its stated intent, when political actors side-step specific 'consent'-driven prohibitions on the collection of personal political opinions by outsourcing their campaigning to intermediaries. While most political parties comply with the prohibition on collecting personal data by asking for consent and limiting trackers on their websites, via the services provided by Facebook and other data brokers data on political opinions come back in through the back door. Though not personally identifiable and thus possibly not in direct violation with GDPR, it is possible to estimate the political opinion of voters without these voters explicitly expressing them. The direct question for policy-makers and data authorities is thus to get statements on the purposes these data brokers use to collect such information, get clear which data sets and tracker profiles are precisely being used, whether and how personal profiles figure into them, and to make sure the links between the micro-targeting components and political parties become more visible. Research in the US has pointed out that when the intentions and data collections behind micro-targeted advertising on for example Facebook is made explicit to targeted users, the advertisement becomes largely ineffective. Because when voters learn about the amount of data collected that delineates them, their desire to be autonomous decision-makers ('self-legislators') instead pushes them to forgo voting for or buying into the ad. (Kim, Barasz, & John, 2018)

In the Netherlands, the potential problems of political micro-targeting are as of yet still relatively mild, and can even deliver positive results. Specifically in the Hague, a small party won the elections and seemed to largely slip through the cracks of our micro-targeting model. When looking at Facebook 'likes', smaller and newer parties seem relatively much more represented, while all political parties are positive about the ability Facebook provides to have more direct and targeted contact with their electorates. In a poll conducted the day after the municipal elections, TOP research found that (far) right-wing conservative newcomer 'FvD' was leading Facebook presence with 148.693 likes, while small (far)-left progressive newcomer 'Partij voor de Dieren' came in second with 120.200 likes. In 2017, small newcomer migrant party 'Denk' came in second, while this year they came in sixth. All three parties, which all three represent ideologically new positions in the parliamentary system, stated that they reach much of their electorate online, and that platforms like Facebook have enabled them to engage with voters much better. While being opaque

and possibly partly illegal, political micro-targeting in the Netherlands has so far thus seemed to have encouraged new entrants in the system, and not hampered new parties from gaining votes.

Economic and Regulatory Incentives for Trackers in the EU

As a final point, for her thesis project, Elsa Rebeca Turcios Rodríguez was able to research the relative effects of regulatory versus economic incentives on the presence of tracker cookies in the EU. What is it that regulators can do now, and where should they target their actions?

To back up this research question with empirical data, we used Open Web Privacy Measurement, a framework developed by Princeton University, to simulate users visiting websites' home page, and we collected data about tracking cookies and cookies notices. We looked up the top 100 country-specific websites for 15 EU countries, and 5 control countries (Australia, Canada, Japan, Switzerland, and The United States), along with 200 top global websites with TLD .com and .org . In addition, we made a cross crawl from the 15 EU countries and the control countries to simulate users' locations for a total crawl of 35,325 websites. We counted 642,362 tracking cookies, 206,787 third party domains, and 217,183 Java Script calls to third-party companies in all websites. In addition, we collected our independent variables from secondary data.

We found that pervasiveness of tracking cookies was 81% in the selected countries not mattering where the users were simulated to be. Besides, top trackers such as Google and Facebook were present in all member states, which implies that a few companies can be encountered on daily basis by users, which should be the target of regulation.

There was nevertheless very large variation in the presence of trackers between EU member states. Websites from different countries present high and significant variability on tracking and cookies notices. The lowest presence of trackers was found for websites based in The Netherlands which have 32.6% less likelihood to have trackers, while the highest presence of trackers was found for websites based in UK which have 3.09 time higher relative risk of having a tracker. Differences in regulation, or how different member state implement GDPR thus matters. But how much?

We observed that whether countries applied consent measures or not did not affect the presence of trackers overall. Also, we observed that countries that have emitted guidance for companies have an increase of 30% in tracking. Moreover, we found that countries that developed fine schemes decrease tracking by 32%, but the overall budget of the data protection authorities did not have an effect. Finally, countries that required to provide more information to users to gain consent to be tracked increased the possibility of having cookies notices on their websites by 6.6 times. However, before jumping to conclusions about whether the differences in the regulatory environment could explain tracking presence across member states, we studied if these differences could be explained by the

role of the market forces and businesses' incentives instead, and to what degree they mattered.

We identified three groups of companies' websites with different business models' incentives that lead to different levels of tracking which we will list from exerting more to less tracking. First, companies which business models' revenue streams are highly dependent on advertisement. Usually, these websites build an audience and give free content to them, and their revenue streams are highly dependent on ads and monetizing their audience. News media belonged mostly to this category, as would Facebook and Google. The second group were companies which business models' revenue streams are slightly dependent on advertisement. These type of business models have other sources of income besides advertisement, and although they track less, they still use ads to have additional income and/or promote their brands. In this group, we found technology and computing, businesses, careers, hobbies and interest, and less so home and garden, Science, education, and food and drink. The third were companies with business models which do not have incentives to use advertisement. Here we found businesses whose main aim is to provide information, and respect the anonymity of users; they are very specific, and do not have incentives to track users. In this group, we found government, illegal content, non-standard content which include adult websites, and health and fitness. Each company chooses for themselves how to bring value to customers, and their revenue streams, and these different business models' incentives to use tracking can lead to market failures, especially for businesses in the first group which profit maximization dependent on ads.

We found that these business models' incentives are powerful predictors of tracking, even more powerful than the different local transpositions of GDPR. This might be explained because the firms are profit maximizers, and some of the business models and their revenue stream have been around for years, and they have been successful, they do not have incentives to change. Due to their self-interest, businesses adapt to the use of tracking at their convenience. Also, businesses have advantages over regulators in terms of personnel, knowledge of technology, and they also produce economic development, so these might be reasons why their incentives become more powerful than the law.

We controlled the provisions of the differences in regulations with business models' incentives and compared them, and some interesting hidden effects from the local regulatory differences arose. We found that consent does have an effect when business incentives are taken into account, and significantly decreases the likelihood of finding tracking by 15%. This might be explained because consent reduces the information asymmetry and Principal-Agent problem between websites and users. Also, we found that countries that developed fine schemes significantly decrease the likelihood of finding tracking by 36%, more than we observe without controlling for business models' incentives. A possible explanation for this result is that fines act as a punishment of companies that do not adhere to the norms reducing business incentives to track. On the other hand, countries that emitted Guidance by Data Protection Authorities did not have the effect intended, but the magnitude of the effect is less. Countries that emitted guidance increase the likelihood of finding tracking by 12%. This result came to surprise us, but we think that might be related

to the fact that guidance might differ from what is stated on the laws or they are more flexible in some countries, so businesses take advantage of this to use tracking. Moreover, the likelihood of finding a banner when more information is required increased to an Incident rate ratio of 8.06. Besides, we noted that EU countries have 24.5% less tracking than our control countries.

These results might suggest that businesses' incentives are powerful and need to be understood and controlled to avoid the called tragedy of the commons, where individuals acting on their own interest deplete a common source, in this case, privacy.

Policy recommendations

Based on our research findings surrounding the Dutch municipal elections, our analysis of how political micro-targeting works, and the research on which incentives matter to the presence of trackers, we make the following policy recommendations, ordered in degree of institutional and political complexity.

1. The first recommendation is that there should be more inquiries into the precise shape of the models used to micro-target persons online. Individual persons have the right to make such inquiries, but very rarely receive the information in a correct, timely and accurate manner. (Verhagen & Andersen, 2018) Companies can meanwhile engage in a host of grey area actions such as 'pseudo-anonymization', where one personal indicator such as a name is replaced by a number. Much better is to opt, as the UK commission on political micro-targeting has done, for demanding fuller openness from data brokers on how they precisely use our information. If and how personal profiles are linked to public data-sets, and what role they play in monetization should concern not only privacy authorities, but also the political parties that liberally hire the services of these players. Next to forcing more openness about the micro-targeted advertising model, our research has shown that giving fines to companies for failing to comply with directives is the most effective way to bring the presence of micro-targeting trackers down the level allowed by law (with 36%).
2. The second recommendation would be to ensure public coordination for online political campaigning. Political parties and regulators could call for an 'ethical pause' regarding political micro-targeting, and sit together to write common 'gedragscode' for political advertising, just as advertising and marketing companies have already done for their sector. Part of such a code could be that 'dark' ads are not allowed, and to cooperate with Facebook in making sure that all political advertisements can be viewed in a central repository. And of course that manipulation is not allowed, and that it is important that personal micro-targeting procedures do not devolve into manipulation. Openness about finances after elections could also be a part of this, mimicking for example the UK's political party expenditure database.
3. Our third recommendation is a bit more far-reaching, and would require long-term commitment. The larger problem surrounding political micro-targeting and the advertising infrastructure behind it is that there is no real public infrastructure for making use of the rights and protections set out in GDPR or in privacy directives. Not only is

compliance with and enforcement of these directives often lacking, which point 1 seeks to remedy. Instead, the current facilitation mechanisms consists of constant consent notifications on websites and in mailboxes, which can be seen as the equivalent of car companies putting up roadblocks in reaction to increased safety measures on publicly financed roads. Through a long and laborious process of requesting personal data from each and every company (the equivalent of asking for car crash test results) can rarely provide meaningful consent. The way out, we suggest, is for governments assist in building the infrastructure that enables meaningful realization of the rights set out in GDPR type regulations.

This could be done assisting in the creation a central 'personal data privacy' infrastructure, such as the 'Solid' and 'Inrupt' projects set out by internet-founder Tim Berners-Lee.[2] In such a system, citizens would have an overview of the different data points collected about their person, and have an overview of the different big data models used to represent to represent their behavior, similar to the public data-points visible in the Central Bureau for Statistics, which were made possible by the previous information revolution. But here, due to the micro-targeting nature of contemporary data collection, control over personal profiles and the 'artificial learning mechanisms' operating upon it would be made explicit. It would be like Facebook, but then so that everyone can see who other than one's friends has access to one's 'profile'.

Within such a new structure for the web, citizens could then authorize different organizations to make use of their data, and possibly even derive money from making their data available to large corporate actors. This more direct approach, which contrasts to a 'tax-and-redistribute' approach advocated in Chapter 4 regarding the digital media-landscape, would empower citizens to be autonomous 'self-legislating' agents in a world of algorithmic governance. In order to maintain meaningful independence from governmental power, governments might restrict themselves to ensuring mere public authorization via DigID of the repository, making the repository itself would held and owned as an independent organization. It could for example be a 'first client' for the 'Inrupt' and 'Solid' system. Security would of course of main concern, but that concern also holds for the current infrastructure, which has not proven to be immune to hacking or abuse.

4. A last, most pro-active recommendation would be for not only public authorities to better facilitate citizens' privacy rights, but for political parties to regain control over their political micro-targeting capabilities away from corporate tech giants. The GDPR is explicitly designed for this purpose, as it makes an exception to the prohibition on the use of personal political data when it comes to organizations of which persons are members. In a digital platform for political parties the creative and positive potential that political micro-targeting technologies can bring to democracy can be brought back under the control of political actors, ensuring better interaction with electorates, enhanced representation and more room for more diverse viewpoints in a democracy

---

[2] https://solid.mit.edu/ , https://www.inrupt.com/

without transferring power to commercial actors. In our final chapter, we will provide a sketch of how such a platform might be designed.

Chapter 4. Political discourse online

In the previous chapter we set out our main findings and recommendations regarding political micro-targeting to provide an answer to our Sub-RQ2, which asked "What particular tracking mechanisms (companies and websites) create an online advertising environment conducive to filter bubbles and echo chambers?" Having shed light on the way the micro-targeting system works, we can now focus on our remaining two central Sub-Research Questions, where Sub-RQ1 asks "In which (echo chamber/filter bubble) directions does online political discourse develop during the process?", while Sub-RQ3 asks "Which algorithms, if any, and which social bots, if any, steer online information consumption in one direction or the other?"

Again, because direct access to the algorithms controlling the Dutch ICT environment was not forthcoming, we had to go about our research in a roundabout manner. Though research points out most political discourse online moves through the Facebook platform, Facebook has made it impossible for researchers to extract any kind of systematic information on it. We therefore focused on two other major platforms: Twitter and Youtube. After a short recap of the potential problem 'echo chambers' or 'filter bubbles' might pose to Dutch democracy, we present our findings on the political discourse on Twitter surrounding the 2018 municipal elections. This research consists of at least three components, and. The first concerns the presence and influence of automated bot accounts on Twitter, which are the best possible proxy for foreign and illegitimate interference. We also specify their relation to particular political parties. The second component of our Twitter data analysis represents the debate online. We show the number of interactions between users affiliated to different political parties, which gives a proxy for the extent to which parties operate in an echo chamber or rather engage strongly with others. Here, perhaps somewhat surprisingly, PVV-affiliated Twitter users are most central to the discourse on Twitter. We have also modeled the topics that played a role in the municipal election and plotted them over time, which for now gives an overview of how the municipal campaign played out online. But this topical analysis, and the more fine-grained 'sentiment' and 'credibility' scores will play more prominently in the third, still to follow component. There, we construct a model that can reveal the civility and credibility of discussions surrounding particular topics in the election and surrounding particular groupings around political parties. This would allow us to say something meaningful about the quality of online discourse, or which topics or parties incite intolerance, versus those that inspire civil discourse.

Next to Twitter, we have also investigated the political direction of content on Youtube, where we discerned a unexpectedly strong bias towards far-right content and against left-wing content. We did not have the resources to leverage a more overarching look at the consumption of online media content by Dutch citizens surrounding the elections, but rest assured in the knowledge that such research and data collection is currently conducted at other university research centers. We will give a brief overview of the results of similar research conducted in the US context surrounding the 2016 elections, which has turned out

to surprisingly re-assuring. Rather than persons being trapped in filter bubbles, a scientific consensus seems to develop that persons online are merely operating in echo chambers, which they do break out of particularly surrounding major political events. We will lastly give a short overview of what fake news is, how Russia has had an interest it, how it can be particularly toxic in combination with micro-targeting and echo chambers, and how it spreads online, before turning to our provisional policy recommendations.

Twitter bot presence surrounding the Dutch municipal elections

The fracturing of news and public discourse in the Dutch, broader European, and North American context is now widely recognized as a problem for democratic deliberation. To the extent that citizens do not agree on basic hard facts and common knowledge, disputes between them are not likely to end well. The point of democratic deliberation, as we conceive it in this context, is not to arrive at consensus, but to reach compromises about how to respond to mutually-recognized political, economic, and policy problems and opportunities in a way that everyone at least understands. When the background of shared common knowledge shrinks or disappears, this becomes untenable. Fracturation in citizens' information sources and networks of trust and reliance thus undermines the capacity of the public to engage in democratic deliberation.

This fracturing is not a new phenomenon. The pillarization (*verzuiling*) of Dutch and Belgian societies lasted approximately one century. However, in the twenty-first century, new media and information technologies have the potential both to increase the amount of fracturation and to screen off members of different communities from one another. Under the pillarization regime, there were only three primary pillars (protestant, catholic, secular). In the current climate, it is possible to have more but smaller groups. In addition, to the extent that these groups share information amongst themselves in closed or semi-private settings, it becomes difficult or even impossible for the media and regulators to engage in fact-checking. Perhaps even more worrisome is that members of other groups are likely to be ignorant of what is presumed to be common knowledge among their political opponents. Under such conditions, instead of opponents engaging in debate, deliberation, and discussion, citizens come to resemble ignorant armies who clash by night.

Moreover, the rise of bots and other forms of automated activity on the Internet may lead citizens to thinking that the views and values of their in-group are much more widely accepted than they actually are. Consider two ways in which this can occur. First, if a citizen's activities online are tracked using cookies, spyware, malware, and the like, it becomes possible to systematically serve that citizen advertisements and other content that confirm what they already think. Dissent and disagreement can become invisible. Second, on social media, bots and other sorts of automated accounts can amplify the messages that a citizen sees or posts, once again making it appear that their views are more common than they actually are.

These are worrisome possibilities. To investigate them, we conducted a study of social media amplification on Twitter. While political discourse no doubt also occurs on other platforms, Twitter is the most open to researchers and therefore serves as our case study.

We set up a streaming API of Twitter to collect all tweets, retweets, replies, and mentions that included hashtags, text strings, and/or handles related to the 2018 municipal election. Monitoring began on 9 February 2018 and stopped 22 March 2018. In total, we collected 13,566,217 items. We then filtered out a large amount of noisy tweets and tweets in languages other than Dutch, leaving us with 795,607 unique tweets for analysis. These represent almost all engagement with the Dutch municipal election that was published by or engaged with at least one of the main parties standing for election in at least one municipality, namely: 50+, CDA, CU, D66, Denk, FVD, GmD, GL, LR, PvdA, PvdD, PVV, SGP, SP, and VVD. We associated with each party its vote-share in the last parliamentary election, as well as its position on a two-dimensional grid in which the x-axis represents economic ideology (socialist to liberal) and the y-axis represents cultural ideology (conservative to progressive). We then associated with each party the official Twitter account representing the party, and recorded both the number of tweets by that account during the relevant period (9 February through 22 March) and the number of followers of that account. In addition, we associated with each party the number of users (whether followers or not) who engaged with the party via use, mention, retweet, and so on, along with the number of such engagements.

While many accounts on Twitter are controlled by human users, many others are partially or highly automated (bots). The distinction is not categorical. For example, an account may be largely manually controlled while also sometimes producing automated activity; by contrast, a heavily-automated account may sometimes be subject to manual control. We used a "bot-or-not" API to deliver estimates of the likelihood of automation for each account in our dataset. Estimates range from 0% (definitely human/manual) to 100% (definitely bot/automated). We treated all accounts at 40% or less as human, all accounts at 60% or more as bots, and all accounts in between as unclear. In addition, we checked which accounts were deleted or deactivated immediately after the election. There are two reasons for this. First, if an account was deleted, then we were unable to estimate the likelihood that it was automated (a bot). Second, deleting an account immediately after the election is a red flag. While it is of course possible to run a legitimate election-specific account, which one deletes when it is no longer relevant, it is also possible that bad actors attempt to hide their illicit activity by deleting the accounts through which that activity proceeded.

Our methodology thus enables us to distinguish four categories of accounts: human-controlled accounts that were not deleted immediately after the election, automated accounts that were not deleted immediately after the election, accounts that may or may not be automated and were not deleted immediately after the election, and accounts that were deleted immediately after the election. This enables us to estimate the

amplification each party received from humans, bots, and deleted accounts during the run-up to the election. Table 1 summarizes our findings at the account level.

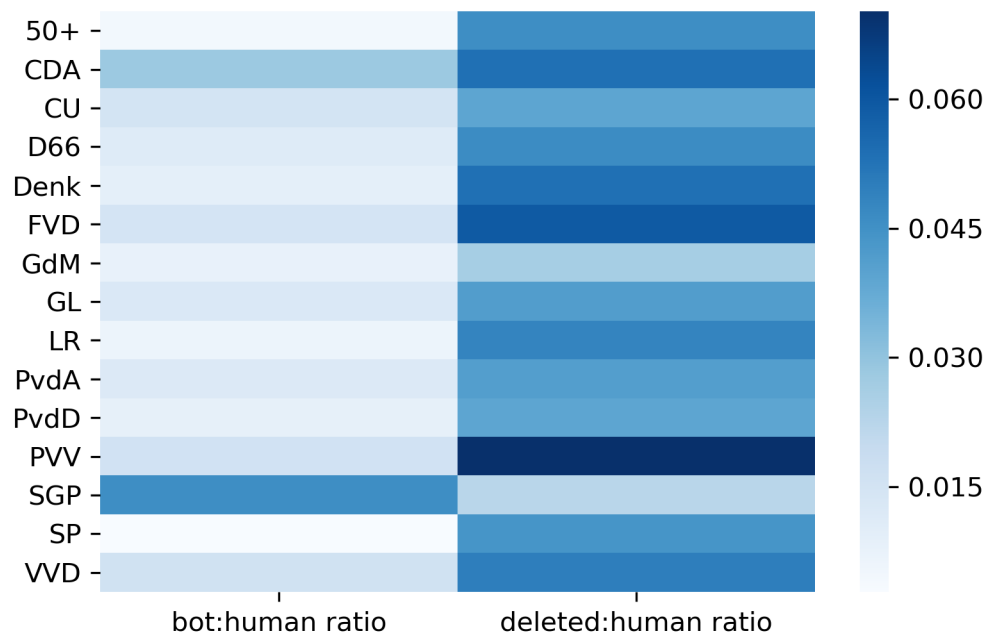**Table 1: human, automated, and deleted-account engagement at the account level**

Table 1 indicates that SGP and CDA had the highest bot to human ratio. What this means is that these parties benefited more than other parties from bot engagement. That of course does not mean that the parties themselves -- or even their Dutch supporters -- were behind the bot engagement. We lack the forensic capacity to determine which actors were behind this activity. Table 1 also indicates that PVV benefited most from engagement by accounts that were deleted immediately after the election, and that FVD, CDA, and Denk also benefited more from such engagement than the other parties.

Automated accounts have the capacity to produce much more content and engagement than manually-controlled accounts. A single actor controlling a farm of automated accounts can mimic the activity of hundreds or even thousands of individual accounts controlled manually. For this reason, it is useful also to examine the data at the level of the tweet in addition to the level of the account. Table 2 summarizes our findings at the tweet level.

Table 2: human, automated, and deleted-account engagement at the tweet level

Party	% tweets by humans	% tweets by bots	% tweets by deleted accounts
50+	85.38%	0.41%	3.09%
CDA	81.08%	1.71%	4.78%
CU	80.09%	1.35%	8.57%
D66	84%	0.75%	4.28%
Denk	82.00%	0.88%	4.95%
FVD	82.62%	0.69%	5.68%
GdM	87.34%	0.69%	2.15%
GL	84.73%	0.88%	4.14%
LR	84.13%	0.64%	4.35%
PvdA	84.26%	0.70%	3.84%
PvdD	86.43%	0.56%	2.99%
PVV	80.37%	1%	4.67%
SGP	74.14%	3.27%	4.83%
SP	86.44%	0.19%	5.58%
VVD	83.39%	0.78%	4.88%

Table 2 indicates that GdM, SP, PvdD, and 50+ received the greatest amount of human-driven engagement. By contrast, less than three quarters of the engagement with SGP was driven by humans. Table 2 also indicates that SGP received the most engagement by bots, followed by CDA and CU. CU received the most engagement by accounts that were deleted immediately after the election, followed by FVD, SP, and Denk. These are all comparative assessments. It's worth noting that, overall, the amount of engagement by both bots and now-deleted accounts was relatively low. Experts estimate that, during the time when these data were collected, between 15% and 66% of all activity on Twitter was driven by bots. In that context, the activity specifically related to the Dutch municipal election seems to have been relatively benign.

Engagement can be further subdivided into types. On Twitter, these including "liking" or "favoriting" a tweet, sharing or "retweeting" it, and replying to it. In general, both favoriting and retweeting indicate agreement or endorsement, whereas replying indicates disagreement. These are of course just rough heuristics, but for data of this magnitude it's impossible to hand-check every form of engagement. With this in mind, we can assess the positive and engagement each of the parties received from humans, bots, and now-

deleted accounts. Tables 3 summarize positive engagement from humans, bots, and deleted accounts.

Tables 3: human, automated, and deleted-account positive engagement

Party	% of retweets from human tweets	% of likes from human tweets
50+	90.15	90.32
CDA	83.77	83.8
CU	82.12	84.73
D66	84.06	86.18
Denk	64.54	64.61
FVD	84.5	86.93
GdM	73.86	78.09
GL	85.05	89.34
LR	79.54	82.44
PvdA	83.95	86.69
PvdD	87.31	91.02
PVV	76.64	79.18
SGP	80.38	72.03
SP	84.54	88.57
VVD	82.83	81.57

Party	% of retweets from bot tweets	% of likes from bot tweets
50+	0	0
CDA	0.44	0.33

Party	% of retweets from deleted accounts	% of likes from deleted accounts
CU	0.29	0.26
D66	0.11	0.11
Denk	0.34	0.35
FVD	0.09	0.07
GdM	0	0
GL	0.08	0.11
LR	0.12	0.04
PvdA	0.14	0.09
PvdD	0.07	0.05
PVV	0.06	0.08
SGP	2.95	5.67
SP	0	0.03
VVD	0.24	0.2

Party	% of retweets from deleted accounts	% of likes from deleted accounts
50+	0.15	0
CDA	1.49	1.97
CU	0.23	0.43
D66	0.64	0.74
Denk	0.42	0.63
FVD	1.28	1.18
GdM	0	0
GL	0.59	0.66

LR	0.4	0.17
PvdA	0.68	0.84
PvdD	0.14	0.05
PVV	1.36	1.56
SGP	0.12	0.15
SP	2.19	1.86
VVD	0.71	1.1

By these more fine-grained metrics, it appears that Denk, PVV, and SGP received the lowest proportion of positive engagement from human-controlled accounts, whereas 50+, PvdA, and GL received the highest proportion of positive engagement from human-controlled accounts. By contrast, SGP received by far the greatest proportion of positive engagement from bots, while SP, PVV, FVD, and CDA received the greatest proportion of positive engagement from now-deleted accounts. Together, these results suggest that the influence of disinformation, which can distort actual citizens' views through bot-generated amplification, is most likely to occur among supporters of CDA, CU, FVD, PVV, SGP, and SP. By contrast, disinformation pushed by bots is least likely to crop up among supporters of 50+, GdM, and PvdA.

Research by Dutch newspaper NRC also revealed that of over 900 Russian Twitter accounts, listed by the US investigation into Russian meddling in the 2016 elections, as attempting to influence the Dutch debate on Twitter in 2016 and 2017. By far, these accounts focused on amplifying anti-Islam sentiment and polarizing attitudes towards migration and refugees. Though not focused around Dutch elections, (a major national election occurred in 2017), the Russian interferences spiked around global Islamic terrorist attacks. More recent reports of Russian influence in the US has shown that over time the focus has somewhat shifted from mere anti-Islam and pro-Trump amplification towards any polarization, where Trump impeachment sentiment as well as leftist fears about prosecution of minorities are also being amplified.

The Russian intent seems to be to effect a decline of trust in shared liberal-democratic institutions like fair and competitive elections, reasonable debate and pluralism, independent media, multiculturalism and basic rights and liberties, among those strong rights guaranteeing freedom of speech. This as part of a larger Russian strategy, where not the absolute betterment of the social, political and economic position of Russian citizens is primarily pursued, but rather the relative decline of Western liberal-democracy. The aim seems to be to bring Western beliefs in free media, in free and equitable markets and in liberal democracy down to the cynical levels that characterizes contemporary Russia. Such

distrust has found fertile ground in the Western imagination since the 2008 financial crisis and its kleptocratic aftermath, the shift of economic power to emerging economies, the collapse of Middle-Eastern states, and the technological transformation of the media. In a sense, Russia arrived at that point decades earlier, after it saw its enthusiastic Post-Soviet embrace of western-style liberalism quickly devolve into kleptocracy and relative decline. Taking advantages of the editorial weakness of the advertising-based model driving new media, it has sought to capitalize on its early retreat from liberal-democratic capitalism, seeking allies for its worldview through polarization among far-right populist parties and other radical, sometime left-wing, fringes. (Snyder, 2018)

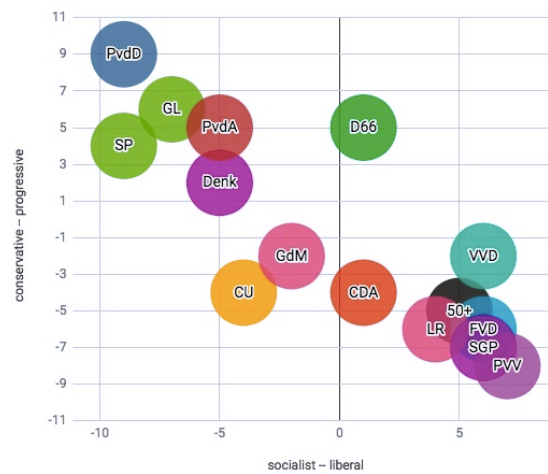## Twitter political discourse direction during the elections

Dutch political parties
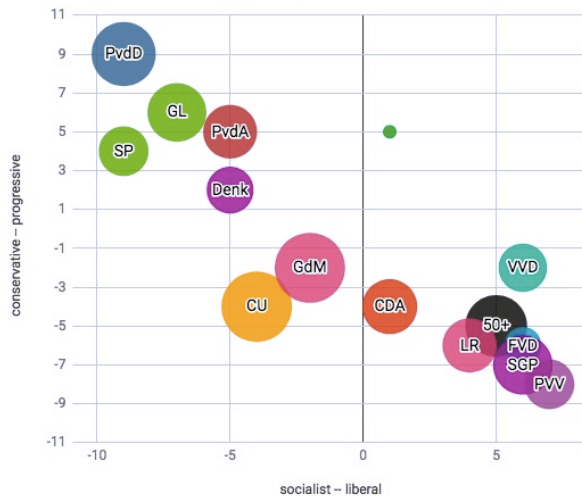size indicates representation in national vote
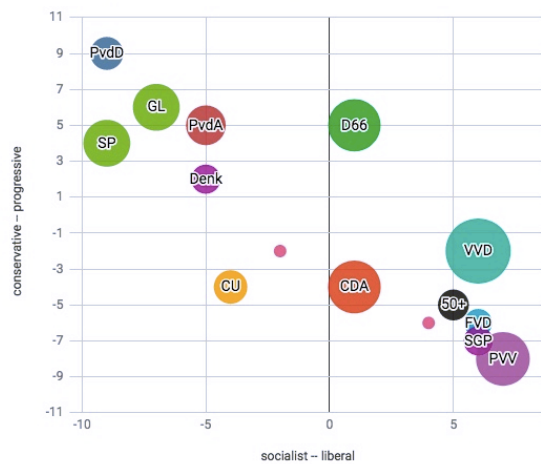
Dutch political parties
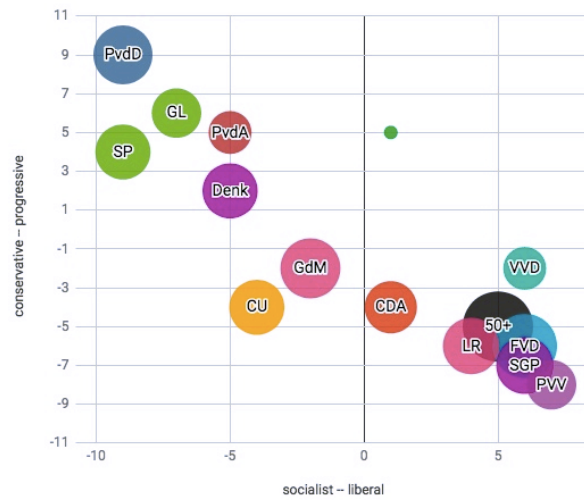size indicates number of Twitter followers

Dutch political parties
size indicates number of tweets by party

Dutch political parties
size indicates number of Twitter users engaging

But is the Dutch political Twitter discourse really so receptive for such anti-democratic Russian influences? Regarding the overall presence of the political parties online, we have seen that in number of Twitter followers a more level playing field exist online than in relation to the party sizes in parliament. This suggests that in terms of followers, all parties seem to get a fair hearing. When we look at the number of tweets by party, the playing field is still relatively even, with some, especially smaller and newer parties tweeting in much larger numbers. In terms of engagement however, a slight bias can be discerned for parties for economically liberal and culturally progressive parties, with outliers for centrist and progressive left-wing newcomers like PvdD, Groep de Mos and Denk. This engagement bias can probably serve as a proxy the larger liberal-conservative bias on Twitter, where topics discussed by these parties get much more attention.

In trying to capture the overall bent of engagement between different party discourses on Twitter, we cast the net a bit wider. We selected accounts that were unofficially affiliated with a party, either through their screenname or their bio, and looked at tweets that linked these accounts to another user that were unofficially affiliated with a party.
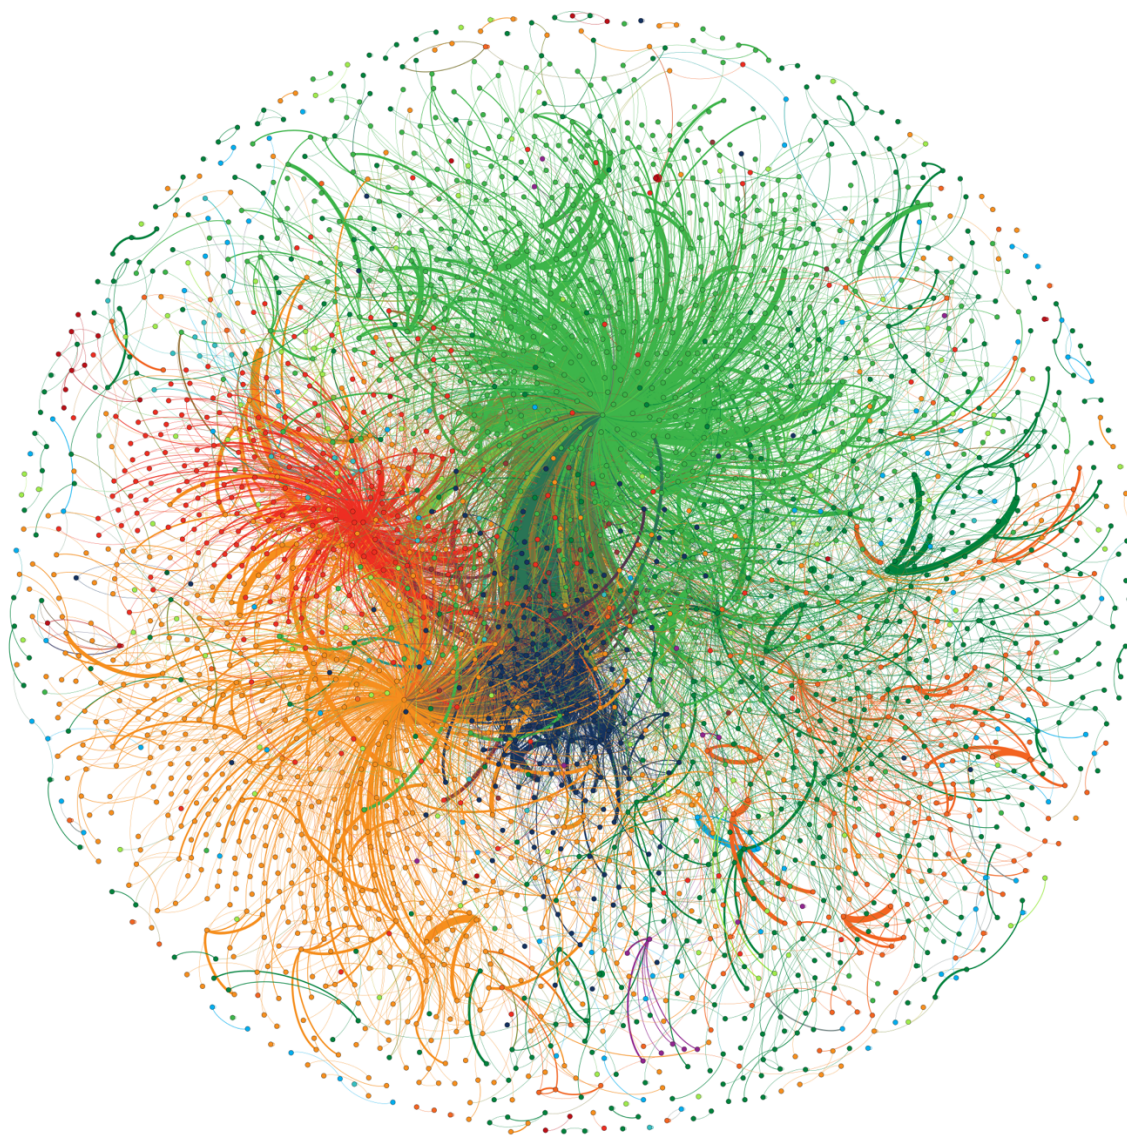
Colour
Party Affiliation

▬ VVD	▬ GL	▬ PvdD	▬ FvD
▬ PVV	▬ SP	▬ 50+	
▬ CDA	▬ PvdA	▬ SGP	
▬ D66	▬ CU	▬ DENK	

What we see here, is that several distinct ideological 'echo-chambers' do exist on Twitter, with politicians and those strongly affiliated to the parties interacting mostly with their peers. The VVD seems to occupy their own space, as do the SP and PvdA, while CDA, D66 and Groenlinks also cluster together in their own space. The PVV also occupies their own space, close to the VVD, but seems central to interactions. PVV affiliated accounts have the most interactions with all other echo chambers, especially the VVD and the centrist and Groenlinks clusters, so much so that the interactions within their own space blur into those

they have with others. This seems to suggest that Russian targeting of those affiliated with 'new right' thought makes strategic sense when seeking to influence the views of all parties. When the central narrative shifts, as it did in the US surrounding calls for impeachment of Trump due to him being potentially compromised by Russia, disinformation campaigns are also likely to shift.

Colour
Party Affiliation

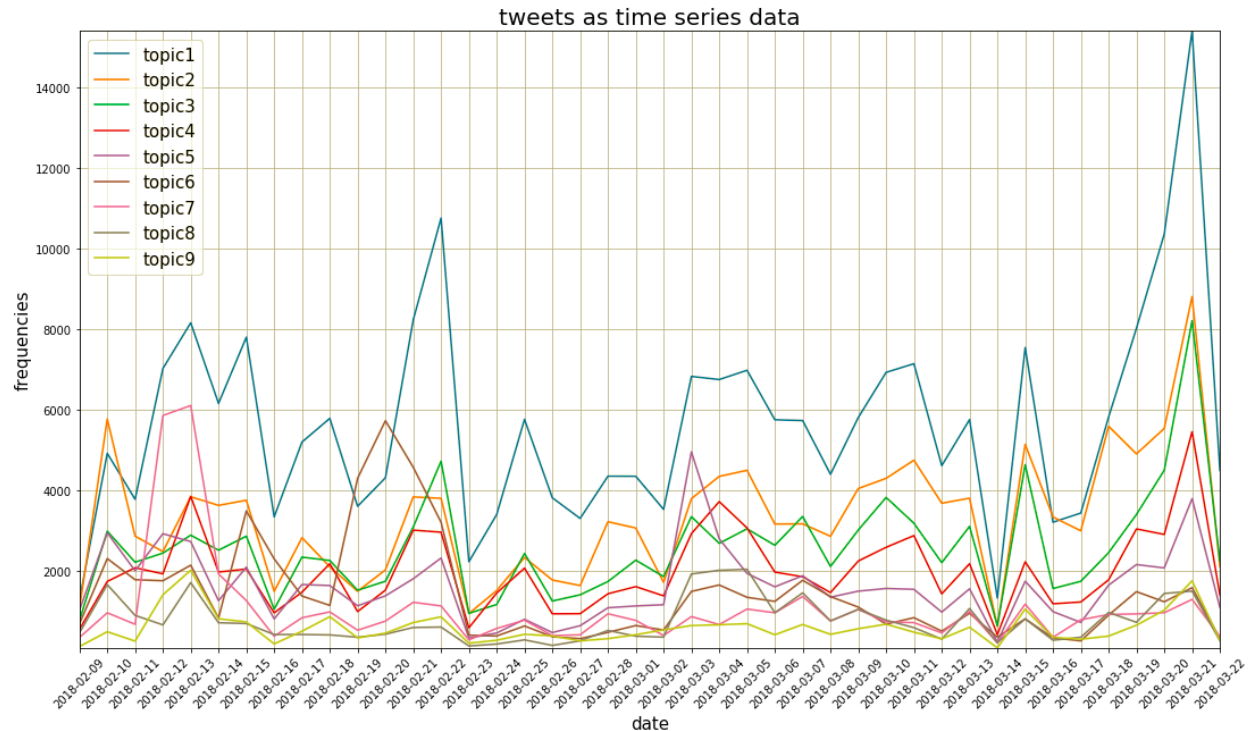|		
VVD	GL	PvdD	FvD
PVV	SP	50+	
CDA	PvdA	SGP	
D66	CU	DENK	
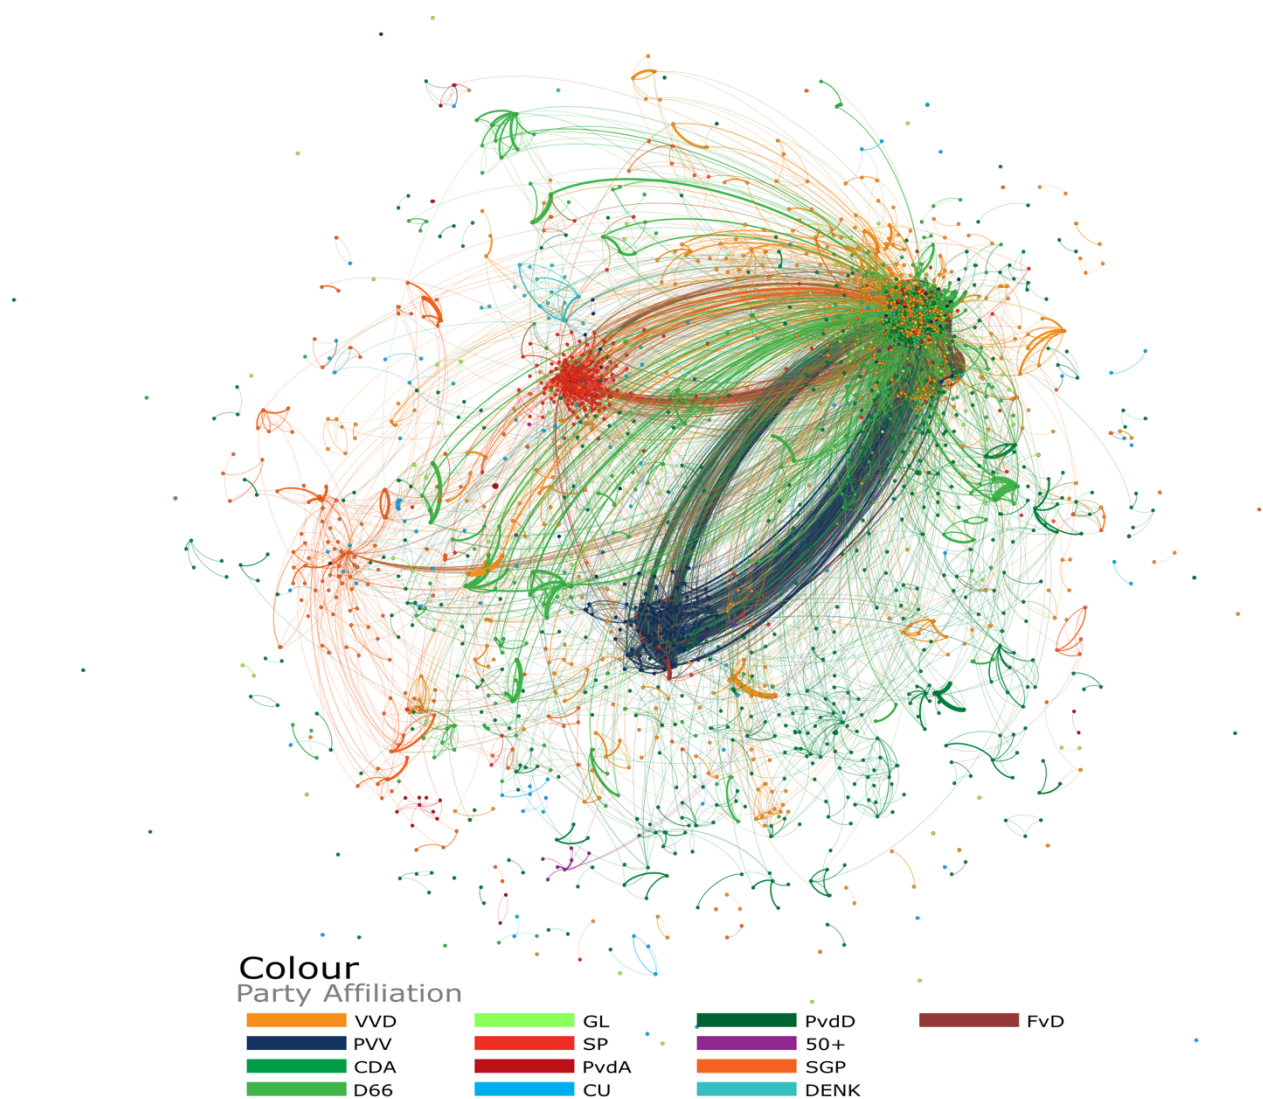
Furthermore, though there is a lot of interaction between almost all distinct ideological 'echo chambers', with PVV most central, the above figure highlights how there is barely any interaction between the PVV and the left-wing cluster of SP and PvdA. If one seeks to speak of a 'filter bubble', those unofficially affiliated with these parties seem to on Twitter be operating within them.

Topic modeling and the quality of discourse

As a stepping stone to our final model on the quality of Twitter discourse, we also analyzed which topics were discussed on Twitter during the municipal elections, and how they developed over time.

Topic 1
vvd weer stem denk wordt verkiezingen echt nederland even referendum

Topic 2
lodewijka meer pvv geert heel pvda moeten wilders hier

Topic 3
kamerlid forum tponl jaar cda vandaag zijn groenlinks rotterdam christenunie

Topic 4
voor mensen ook goed lijst wel worden ers gaat waar zegt dan politiek

Topic 5
amsterdam grootste hoe grote wel islam fvd doet kan heb waarom gewoon moet waarom

Topic 6
kajsa baudet uit eqexyokei ollongren weg minpres van dus thierry ben togetter hebben

Topic 7
zijlstra kamer mee doen lijsttrekker debat campagne maken halbe sgp daar cu staat tijd komt kaag

Topic 8
nog aan ananninga maakt toch partij eens vindt penthouse ddstandaard pechtold joostniemoller

Topic 9
politieke mark mag net zeggen non heeft rutte alleen gaan maart poetin stemmen minister

tweets as time series data

There were some topics we expected to see, such as the 'racism' debate surrounding FvD in Amsterdam in Topic 5, which peaked early on, the debate on fake news on which Baudet and Ollongren prominently positioned themselves in Topic 6, prime minister Rutte in topic 9, D66 leader Pechtold in topic 8 and his failure to register his gifted 'penthouse' apartment, as well as the scandal leading to the resignation of Foreign Affairs minister Zijlstra in topic 7. Topics such Topic 1 on the VVD and Denk we didn't expect, but likely focused around the unwillingness of VVD politicians to debate Denk politicians in Deventer. Topics 2 and 3, on the PVV versus the PvdA and on CDA versus Groenlinks were less clear, as was topic 4.

Going forward, we will combine this topic modelling with the sentiment analysis of our tweets and with the party and ideological clusters, in order to see which topics and which parties inspired particularly positive or negative discussions. We will also try to outline the 'epistemic' positions of each topic discussion and ideological party cluster, in order to reveal how well-informed those engaging in the debate are. This will be measured in terms of the diversity of each' accounts networks, and its (hearsay) distance from primacy sources.

Platform bias and 'Fake News'

We have also analyzed how the Youtube recommendation algorithm functions by automating suggestions, based on a representative input of political party search terms. By collecting video recommendations based on the search terms 'VVD'. 'PVV', 'PvdA' and all other nationally represented political parties, and weighing these representations according to

the number of seats they have in Dutch Parliament, we were able to obtain a list of recommended political Youtube videos that is representative for an undecided, or politically neutral, voter. That (shortened) list of 1000 video recommendations consisted of about 58% neutral content, mainly those linking to traditional media sources, of about 33% either pro-PVV or pro-ForumvoorDemocratie videos, and about 9% from left-wing and assorted remaining sources. When searching for political content on Youtube, one is three times as likely to encounter 'populist far-right' content than that of any other political ideology. Furthermore, when for example searching for content from a left-wing party like 'Groenlinks' zero results are recommended that present the party in a positive light, while for almost every party content is recommended that comes from FvD accounts or accounts presenting PVV and FvD in a favorable light.
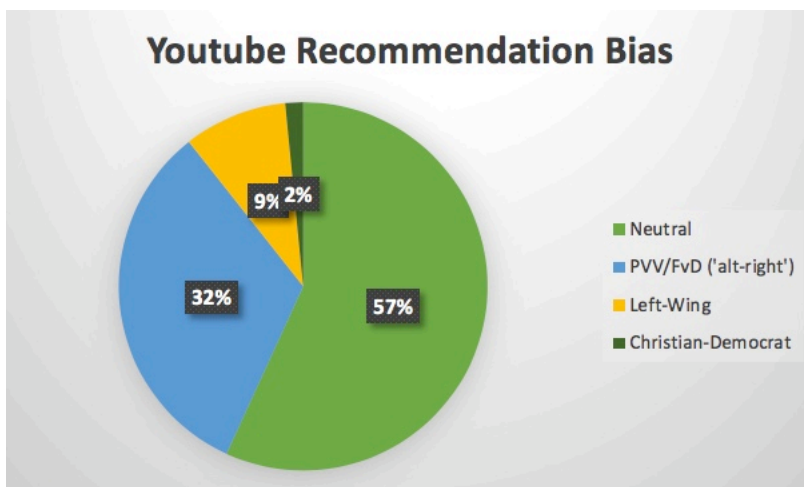
Regarding Twitter, newspaper the Volkskrant found that for new Twitter profiles, the suggestions offered by the Twitter algorithm when expressing one's interest in 'Politics' lean almost exclusively right-wing populist. When asked for explanation, Twitter responded that this reflected the state of Dutch twitter discourse; this is partly confirmed by number of interactions by PVV supporters on Twitter. Nonetheless, even as this would vindicate the 'neutrality' of the algorithm in merely reflecting actual discourse back, a right-wing populist bias is still present for those citizens entering the platform without preferences, particularly on Youtube.

This algorithmic bias also likely drives the strong bias on Youtube, and can be added to the core weaknesses of the ICT environment that poses a threat to democracy, as it reinforces bias to new users. Instead of suggesting mostly right-wing Twitter users, or recommending only videos from one ideological end of the spectrum, algorithms that select content for users could present accounts based on the diverse array of political clusters very much present on Twitter and less visibly on Youtube, and sort them either as a level playing field or as party vote share in parliament. Platforms might seek to emulate the Dutch 'omroepen' system, where distinct ideological views are represented by different media organizations.

Fortunately, this could be easily done technologically, without the need to ministerial interference in public broadcasting that characterizes the 'omroep' system. One condition, now as much as going forward, would have to again be that platforms open their selection algorithms up to public scrutiny, just as the free press does regarding their editorial practices.

So how does the spread of disinformation, false news or 'fake news', that has not been shown to be explicitly Russian controlled, relate to algorithmic bias, the centrality of PVV-type discourses, and the quite balanced presence of different ideological 'echo chambers' on Twitter? Though a conclusive answer to this question will have to wait until our fuller analysis of the quality of Twitter discourse has been completed, several things can be said.

The first is that it is extremely difficult to pinpoint what is fake or false news, and that doing so puts researchers and politicians on a dangerous slippery slope. A couple of research projects have so far been successful (Guess, Nyhan, & Reifler, 2018)(Allcott & Gentzkow, 2017; Vosoughi, Roy, & Aral, 2018). In these analyses, fake news surrounding the 2016 US elections is narrowly defined through the reports given by fact-checking websites. In the Netherlands, professional fact-checkers have not flagged the consistent presence of false political reports. Other research projects however, such as that carried out by the Oxford Computational Propaganda Project (Vidya Narayanan, Vlad Barash, John Kelly, Bence Kollanyi, Lisa-Maria Neudert, 2018) dangerously miss the mark in its selection of what counts as fake news or not. Among the websites listed as 'fake news' were not only new 'far-right populist' news platforms such as Breitbart, but also conservative journals with long and strong editorial histories such as the National Review, and respected polling firms like Rasmussen Reports. In the Netherlands, the helpful repository 'Hoax Wiki', and in the EU the EU vs Disinfo body, has also often too hastily tagged journalistic content as 'fake news'. This biased listings in turn only reproduces itself in biased conclusions, which themselves damage shared trust in not only media but also academic institutions.

The conclusions of studies on fake news surrounding the 2016 elections that did use reliable and unbiased lists of fake news reports pointed out that fake news website production and consumption was overwhelmingly pro-Trump. Moreover, the echo chambers through which such false reports were consumed were reported to be quite 'deep', with high numbers of false articles being consumed per user, but also quite 'narrow', so that the groups consuming fake news represented only small >10% fragments of the public. (Guess et al., 2018)

Research also showed that social media websites play a strong amplifying role in fake news consumption. Where around the 2016 elections visits to regular top news websites was directed from social media at around 10%, over 40% of traffic to fake news websites came from social media platforms.(Allcott & Gentzkow, 2017) Facebook plays an outsized role in directing this traffic. Research on overall media consumption online carried out in the US also revealed that there is a surprising *lack* of divergence in media consumption between Republicans and Democratic. (Guess, 2018) Users check similar centrist news sources often, and even those caught in deep 'echo chamber' media environments refer back to centrist

news websites when major political events occur. Whether such 'fake news', Russian interference, and micro-targeting vulnerability swayed the elections remains a contentious point.(Jamieson, 2018)[3]

Further comprehensive research was conducted on the spread of false news on Twitter by Vosoughi et. al. in Science (Vosoughi et al., 2018). By looking at Twitter data covering over 10 years and affecting 3 million people, the researchers showed that false news not only reaches far more people than true facts, but also that it spreads much faster. They found that

> "Falsehood diffused significantly farther, faster, deeper, and more broadly than the truth in all categories of information, and the effects were more pronounced for false political news than for false news about terrorism, natural disasters, science, urban legends, or financial information. […]False news was more novel than true news, which suggests that people were more likely to share novel information. […] Contrary to conventional wisdom, robots accelerated the spread of true and false news at the same rate, implying that false news spreads more than the truth because humans, not robots, are more likely to spread it." (Vosoughi et al., 2018)

Most interestingly, 'influencers' and central nodes did not play a large role in the spread of false news, but it was rather the larger number of less-linked accounts that spread false news for the sake of novelty. A worrying lack of journalistic ethics among internet users thus drives false news, where sources are not checked and hearsay drives content. The 'seeing' and 'reacting' elements that drive the advertisement-based online attention economy thus overtake the 'reflecting' part, disfiguring, as Urbinati would say, democracy in a reactive populist and complacent audience democracy manner, pushing truth out with falsity. The desire of governments and politicians to curb such disfigurements of the freedom of speech online is understandable, but might, when not done carefully, further undermine trust in liberal-democratic institutions.

Policy Recommendations

Our research on Dutch online political media around the 2018 municipal elections revealed both the positive potential of new media and their weaknesses. On Youtube, Twitter, and probably on Facebook, algorithms present an ideologically biased selection of content, over and above the perceived overall slant of online discourse, which tend in a conservative right-wing direction with PVV nodes appearing central. Digital micro-blogging platforms also create an environment conducive to deep echo chambers and the rapid spread of fake news, even as the Netherlands has so far proved quite resilient. Reports of false news have been very scarce, but Russian accounts have attempted and are likely to still influence discourse online. Whichever narrative is most strategically central online is also the one prone to most disinformation influences. Any citizen online should be safeguarded from nefarious interference, in order to ensure a fair contest of ideas that drives truth-finding

---

[3] See: https://www.newyorker.com/magazine/2018/10/01/how-russia-helped-to-swing-the-election-for-trump

in a democracy. If not, the collective process of fact-finding and opinion-formation can spin out of control. In the US, a growing part of the electorate now seem eager to throw the media baby out with the 'fake news' bathwater, as 43% of Republicans in the US now believe that the president should be allowed to shut down media. (Fredericks, 2018)

We also saw that diverse viewpoints were present on Twitter. Moreover, new political parties seem to benefit greatly from all this. Citizens can quickly align themselves with new actors, engage in the discussions driving these movements, and engage critically and equally with politicians and newsmakers. The promising new digital democratic world that was expected to be born does still shine its light through cracks in the digital corporate firmament.

Two classically competing principles of the liberal representative democracy thus need to be balanced anew online: an almost unlimited right to *freedom of speech*, which is required for individuals to express their thoughts, beliefs and interests and for societies to approximate truth, and the need for *representation*, which is required to manage the scalar and temporal complexities that deliberation between very large numbers of people poses. Where ICT media has positively supercharged freedom of speech, it has unhinged traditional representative nodes. As Arendt saw, such changes in the collective ordering of representation will affect shared notions of truth and opinion, and can dangerously destabilize politics.

John Stuart Mill already saw clearly how to balance these competing concerns. For him, his radical defense of freedom of speech and the extension of the right to vote to women had to be balanced with 'plural voting', where those with more education would wield more votes than those without. When one side, of individual freedom, increased participation, and greater inclusion is enhanced, so should the means by which it is ensured that the public at large can in turn make use of only the best quality facts and opinions.

'Plural voting' is of course not a measure fit for our time. Since Mill's time, several other intermediaries have come to play the role of such representative agents; primary among them political parties, academic and educational institutions and media organizations. We see that in places where fresh political parties can easily enter parliament and where media and educators are trusted highly such as the Netherlands, these representative agents are able to help bring out the positive aspects of the digital revolution. In places with very rigid two-party structures, such as the US and the UK, and where public funding for and trust in media and education has been steadily declining for decades, the expansion of the freedom of speech that accompanies the technological transformation of the public sphere shakes society much more violently.

In order to enhance the quality of representation in an (online) society where strong freedom of speech protections and increased participation remain of the highest importance, the Dutch system thus already has several strong advantages over other societies. Being a multi-party system, increased personalization can help the quick rise of new political movements, which can then easily enter local and national parliaments. The required coalition-building afterwards ensures an increased plurality of perspectives might enable fruitful coalition-building between diverse parties in the Hague and other

municipalities revealed. High citizen trust in neutral media arbitrators such as the news-organization NOS, and the successful emergence of centrist online news-providers such as NU.nl have together with the plural representation in the 'omroep' system of different viewpoints that align with both new and traditional political ideologies ensured that the representative agents in Dutch society have been well-equipped to deal with the digital revolution.

Trust in the institutions of liberal-democracy is in decline all over the world (Economist Intelligence Unit, 2017; Foa & Mounk, 2017; Varieties of Democracy Institute, 2018), and media autonomy, freedom of expression and alternative sources of information being among the democracy metrics that have undergone the greatest global decline in recent years. In the Netherlands, trust in the media is very high, while citizen concern regarding disinformation is extremely low, and the percentage of citizens that have reported to come into contact with disinformation is with 10% the second lowest globally (Commissariaat voor de Media, 2018; de Cock Buning, Eljon, & Bune, 2018; Fletcher, 2018). Hastily government censure of 'fake news', which the French government has instated for election periods, or even making platforms liable for what is and is not 'correct speech' is thus in no case required here.

Instead, we would recommend further strengthening Dutch 'representative agents', so that the positive 'inclusion' and 'freedom of speech' aspects of the digital media revolution can further enhance liberal democracy. This means more attention in education on technology and media and further pushing digital platforms to take up the traditional responsibilities that media companies and publishers have. It also means pushing television and print media to increase their content online, while encouraging political parties and parliament to make fuller use digital means through a political party platform. The precise shape of that last recommendation will be fully outline in the final chapter, while we will shortly elaborate on the others here.

1. Better citizen education on journalistic practices should accompany the inclusion of citizens in media production online. Fortunately, on the 5th of March the Dutch NOS did just that, by hosting a three-hour broadcast on the meaning and influence of fake news, interspersed with short lectures from journalism schools. What gets taught in journalism schools, namely how framing, narrative and tone are all used to color any news in one direction or another, should become standard practice in secondary education. This would tailor well with the experiences of teenagers, who already engage massively in media production through social media. Examples of personal social media experiences and potential abuses could accompany larger issues on the role of truth and facts in a democracy and the shared duties of citizenship.
Additionally, journalists, politicians and academics, who do often act as 'gate-keepers' online could in their interactions with 'followers' present not just stated facts or considered opinions, but also outline and include others in the process and editorial practices that have let to said facts or opinions. In order to, together with

digitally enhanced political parties, ensure that citizen participation can through emancipation also lead to more meaningful power-sharing.

2. During our research period, the scandals affecting digital media platforms also rightly led to some meaningful changes in the editorial policies of big technological corporations. Twitter has since May started blocking over 35 million accounts every month, eradicating the swaths of fake accounts plaguing the platform. This is expected to significantly hurt their profits, as their advertising revenues is based on numbers of user accounts, and shows commitment to editorial responsibility. Facebook has also significantly changed course away from profit maximization through attention maximization, or the quantity of time spent on the platform, and towards 'time well spent', or the quality of interactions one has on it. This has greatly impacted their revenue, with Facebook shares declining 20%, effecting a record 129 billion loss. The company has also attempted to get rid of election meddling, by making ad buyers provide a postal address in the country they are advertising in, and has started setting up procedures together with 'Social Science One' to allow researchers to analyze its data. We cannot but be encouraging about this long overdue course correction.

   Both companies, as well as Youtube, should however go much further. Though it might sadly clash with the creative freedoms the internet has so far licensed, further tracing online identities to actual persons should continue. Newspapers and other publishers also always check and retain the identity of their sources and contributors, as it is essential for democratic citizenship that one's personal expression is coupled with the responsibility to own up to one's views. Exceptions, such as in the case of whistle-blowers can be made, and several partial proxies could online serve to identify persons, while assuming multiple identities online should still be possible. A better personal data infrastructure, such as that proposed through 'Inrupt' by Berners-Lee, could facilitate this change.

   Platforms should also open up the choices they make in algorithmic selection to other newsmakers, to academic researchers, and potentially to the general public. These algorithms should also give a more plural representation of viewpoints; within one's personalized 'stream' opposing viewpoints can be presented, selected from the large number of different posts currently often hidden by algorithmic selection. Presenting both sides of a story is a basic journalistic responsibility, and without it, it is almost impossible to present meaningful democratic deliberation as a contest of ideas. As philosopher John Rawls argued, in order to accommodate the 'basic fact of pluralism' in liberal democracy, different 'freestanding political conceptions of justice' should reasonably find ways to ensure an 'overlapping consensus'. But when we do not whether or how our views differ or are in agreement, we lose our capacity to govern justly together. (Rawls, 2005)

3. In addition to enhanced responsibilities for online platforms, it is essential that media organizations with high journalistic standards have the resources to extend their presence online. In the Netherlands, public broadcasting budgets are being continuously cut because less and less people watch television or listen to radio as

they move online. Print media has fortunately already changed much of their approach. After long periods of declining readership, newspapers have made their content not only available online but often in a way that is more attractive than on paper. For TV this has yet to happen, even as radio has caught on to the podcasting trend. Television stations might cooperate more closely with Youtube, Facebook and others, and governments could ensure these changes are accompanied by the right financial incentives. Offline, publishers also pay writers and content creators, and while Youtube has spawned a whole host of well-earning vloggers, the stream of money from platforms to high-quality journalism is lacking. This might be done by taxing the platforms, in line with the EU's proposed 'link tax' or by ensuring that on a basis other than mere viewership, Youtube (Google) and Facebook start to remunerate quality journalists for the production of their content. As our research showed, the 55% of political content on Youtube that is neutral comes almost exclusively from public broadcasting sources, while the revenue made on it is currently diverted away from their source towards often anonymous uploaders.

Intermezzo: Digital platforms for Political Parties

Next to the more 'reactive' policy recommendations that have been advanced, it is also possible to engage more constructively with the relation between Dutch democracy and the digital environment. The facilitation of digital rights through a public infrastructure was one way to do so, but the number of possible 'responsible innovations' are much more numerous.

In a multiparty democracy such as the Netherlands, the usual route through which such constructive changes are put forward is the party system, where political parties themselves propose new responsible innovations that can strengthen democracy. In our research, we have been positively surprised by how receptive political parties have been to suggestions. But we are also aware that time and resources are scarce in day-to-day politics, so that aspirations such as these often get pushed to the background. Because of the special role of the State Commission, which stands apart from direct party political competition, and because we think the most urgent responsible digital innovation concerns the infrastructure of political parties and of parliament itself, we have put some thought into how the design of a digital political party platform might look like. The need for internal party democracy in each and every political party in a liberal democracy is here presumed.

It is important to focus on increasing digital participation in political parties because parties serve as essential 'representative intermediaries' whose membership has steadily declined, even as their number has proliferated. In political science, the larger crisis of democracy is often attributed to this declining party membership, and the concurrent rejuvenation of democracy with the revitalization of the party system. In our view, the digital revolution and the attendant increase of digital participation in online platforms and social media can and should be understood not just as a change in the way we consume and produce media, or the way in which political parties can advertise and communicate their views, but also as a fundamental change in how citizens participate in collective decision-making. Political parties have traditionally functioned as the participatory avenue that accompanies the opinion-formation and fact-finding functions served by the media, and a decline in party membership thus also signifies a decline in opportunities for political participation. Linking increased but corporately misdirected digital participation to declining offline party participation thus presents an unique opportunity for digitally enhancing participation in political party structures.

A responsible digital innovation for political parties can primarily help to put politicians on a more equal basis with members, bringing 'the parliament' closer to 'the people'. A platform might achieve this by opening up the ordering of pre-election party lists to a more natural and fluid digital deliberative process, as well as by having elected politicians develop a stronger bond with the members through constant digital member consultation. A digital platform for political parties should thus consist of two elements: an online deliberation environment, and an online consultation environment. These are kept distinct, as for consultation members should be able to cast their vote, and thus retain a measure of a singular and verified identity, while for a deliberation environment the largest array of

possible 'half-baked' ideas and proposals should be put forward, which can be encouraged by allowing every member to anonymously post thoughts and comments without verification. Only for those seeking to rise in popularity among the ranks in order to compete for a place of election party lists would the need to identify oneself with one's ideas become very attractive. In addition to the expressive 'deliberation' and the reactive 'consultation' environments there would also be a 'master narrative' environment, which forces a reflective dimension on the free exchange of ideas and the day-to-day reactions of parliamentary member consultation. Let us highlight these environments in turn.

1. An online discussion environment would serve to allow elected politicians and all members to present, discuss and assess any kind of idea on an equal basis. As with for example Facebook or Twitter, ideas can be proposed and everyone can 'like'. 'dislike', or comment in such a 'massive online open deliberative environment'. Algorithms could assist in ordering the ideas into structured logical arguments, with platforms such as Kialo already providing technological backbone for this.

The 'idea ranking score' of users on the platform would furthermore be indicative for a place on the electoral list; elected and aspiring politicians would thereby be encouraged to build up a substantive relationship with their constituencies. Not only likes but also dislikes could count in final social ranking scores, so that the experimental character of democracy is stimulated. Many of these 'massive open online deliberative' environments are already available, and could be prototyped in a relatively short time.

2. An online consultation environment would secondly ensure closer interaction between represented politicians and party members. Here, elected politicians should be able to constantly put their parliamentary contributions to voting, motions, amendments and initiatives up for consultation with members. This could for example be done in a short film in which elected politicians explain what the topic at hand is, and what their proposal for legislative change would consist of. The legislative text can then be opened up in a 'Wikipedia' structure for changes by members, in which every change can be voted on, while members can get 'notifications' about decision moments. Algorithms could here again assist in translating proposed changes into legislatively appropriate language. The parliamentary representative would after a consultation period then promise to take this input with him for the final vote, or possibly bind itself to defend the member amendments. Each member would have a vote for the largest possible number of legislative actions of the party, and members can choose to use these directly or delegate to them another member. By default all such votes would be delegated to elected politicians (often party leaders). This means that at first, nothing would have to change without members explicitly wanting to participate. Members would also be able divide their votes among various other politicians, or delegate them to other members with certain expertise.

Together with the 'social ranking score' from the debating environment, these delegated voices form a basis for the natural creation of electoral lists and regularized interaction with members, so that politicians and members are given the opportunity to profile themselves online in the party on an equal basis. These 'liquid democracy' platforms are also available,

for example as used by European Pirate parties, the Five Star Movement in Italy, and by Wikipedia and other freely available democratic software platforms.

3. In addition, collective 'election narrative' creation could serve as a master narrative that gives direction to the consultation environment, so that it is more difficult for both members and politicians to deviate from the agreed election program. This would enhance the representative aspect of democratic politics, balancing out the increases in participation the platform would bring with space for reflection and vision. After the narrative 'election program' has been agreed upon, the relevant paragraph from the election program could then appear next to the proposed legal text in the consultation environment. Additional voting thresholds could make it more difficult for both members and election politicians to deviate from the 'election program' in day-to-day legislation, so that a measure of long-term party discipline could be maintained.

The 'election narrative' text would however also be open to changes by members, but at other regularized intervals, prior to a parliamentary seating period. This additional restriction could in the Dutch system be mirrored in a 'coalition agreement' from which members' consultation could not deviate when their party participates in a government. This, because it should be very difficult for both members and politicians to deviate from coalition agreements as well to retain stable executive coalitions.

For many important elements it is possible to build in flexibility and only after a long test period to reach a final decision. For example, how many members should participate in a consultation or discussion to give this weight or to obligate politicians, which majorities are needed (50%, 75% etc), how many voting rounds are needed, or to what extent participating members should be demographically similar to the diversity within the party as a whole. The platform should ideally be introduced in phases, where after a development period (until for example the next elections), first the electoral list would be partly determined on the platform, then a test period of four years for the consultation platform would take place, after which the election program master narrative could be opened up to wide member participation. In this way, not only could members become increasingly be more involved in parliamentary politics and legislative power, but they could also learn about the intricacies of democratic decision-making.

Chapter 5. Conclusions

So what are the conclusions that can be drawn from the research? After recapping the sub-research questions and their main conclusions, we will be able to give an answer to the main research question, which was whether the ICT environment has delegitimized the 2018 municipal elections. Having answered the main research question, we will then shortly recap our policy recommendation on both what action should be taken and which actions are likely to be counterproductive.

In chapter 4 the topic of the first Research Questions was treated, which asked in which 'echo chamber' or 'filter bubble' directions the online political discourse develop would during the municipal election campaign. We modelled the main elections topics of the on Twitter, and found that 'fake news' itself became one of the central election topics. Moreover, we found that PVV affiliated users were the most interlinked users on the platform. Other political groupings, such as the liberal VVD and the centrist CDA, D66 and partly Groenlinks, each isolated itself more in 'echo chambers', though they interacted quite a lot with the central PVV-users. The left-wing block of PvdA and SP users did interact somewhat with non-PVV groupings, but most closely resembled a digital 'filter bubble'. The research was however not fully conclusive, as interaction outside the Twitter platform is likely to have also taken place, where a decrease of fragmentation has been documented in other countries. We also saw a relatively level playing field between the parties on Twitter in terms of followers and posts, though there was a slight skewing towards all conservative-liberal parties in terms of overall interaction. This can probably be ascribed to the slight ideological bias in that direction that Twitter has said 'neutrally' reflects their user content.

The topic of the second research question regarding which particular tracking mechanisms, companies and websites create an online advertising environment conducive to filter bubbles and echo chambers was answered mainly in Chapter 3. There, we found that Facebook and Google were the main trackers present, but that a whole host of smaller companies offered targeted micro-targeting capabilities. Political parties did not make use of these political trackers on their own websites, but reported to have used the services of political micro-targeting companies and the capabilities offered by Facebook. Three elements were identified as essential to political micro-targeting, which were firstly the advertising infrastructure provided by large corporations like Facebook or Google, secondly the presence of individual trackers and cookies, and thirdly the big data analysis capabilities required to link voting and demographic data to the Facebook infrastructure and individual tracking profiles. We ran a successful initial big data analysis for the city of the Hague, showing that it is possible to predict voting behavior based on demographic data. Though we did not make use of individual tracking profiles, it seems likely that the creative force of local party politics in the Hague was able to overcome increased political micro-targeting capabilities. In our big data analyses the election the main winner 'Groep de Mos' was not revealed to have a target demographic, and would thus have targeted its constituency outside of the online environment, or with mere brute all-out advertising force.

We also find it safe to assume that increased micro-targeting capabilities and the personalized online infrastructure are likely to contribute to party fragmentation, as all national newcomer parties benefit most from the online ability to reach ever smaller segments of voters. But specifically in the Hague, 'Groep de Mos' was able to gobble up the electorate of many such smaller sub-group parties to become the largest party. In local politics, it appears that ample room exists for political craftmanship to counter fragmentation. Whether this also holds for the national level remains to be seen, as the larger electorate and further distance from citizen's direct affairs might still be open to further differentiation. Such digitally-enhanced 'fragmentation' might not necessarily be bad for the Dutch parliamentary system, as long as it is balanced with a post-election willingness for coalition-building and consensus among parties. But it could be potentially strengthened with more direct voter involvement in coalition-building. The recommendation of the State Commission to let citizens directly elect a 'formateur' during national elections can be embraced in this regard.

Regarding the third research question, which asked which algorithms or social bots steered online information consumption in one direction of another, we found a low presence of social bots, which engaged slightly more with strongly conservative-liberal parties on the political spectrum. Research into the content of Dutch Twitter bots confirmed this, finding that anti-Islam rhetoric with the intent to polarize was its main direction. Algorithms on Twitter and particularly on Youtube also presented a far-right-wing bias. On Youtube, video recommendation almost never presented left-wing viewpoints, and suggested conservative right-wing content in one in three cases.

For the fourth research question, which asked which themes, political parties and politicians in the Hague and the Netherlands more broadly characterized the municipal elections, we found that these focused a lot on the expected and often centrist themes such as housing, security and migration. As said, de Mos surprisingly won in the Hague, with an eclectic mix of ideological positions. On the national level the discussion on the role of the internet on democracy itself became a major theme, when the government expressed its support for an EU-led initiative that incorrectly labeled Dutch news as disinformation. The more surprising theme of 'racism' played a role early on, but was diffused after the first election debate.

What does the ICT environment do to democracy?

Over and above the specific research question, the research delved deeper into two specific topics; ICT media and political micro-targeting. From the perspective of democracy, these two elements can be seen as two digital incarnations of the two elements of the democratic rechtsstaat, the majoritarian component in online ICT deliberations and the individual right-protecting aspect in micro-targeting.

We found that political micro-targeting primarily poses risks for the protection of rights, autonomy and 'self-legislation', and that the precise functioning of the models driving micro-targeting should therefore be open to the public. This could be facilitated by a public digital platform for personal digital rights. The new ICT media was meanwhile found to be

subject to similar weaknesses that majoritarian institutions can pose, where a balance need to be struck between representation and participation. While supercharged possibilities for freedom of expression and participation in collective truth-finding and opinion-formation are facilitated by ICT media, the capacity for representation can be undermined through the weakening of intermediary institutions such as editorial journalism, journalistic ethics and political party participation. In international research, we found that these intermediary institutions are especially vulnerable online, and have been exploited through the Russian 'fake news' attacks. The Netherlands proved quite resilient, educating citizens on the changing media environment, enabling a diverse contest of ideas between differing opinions, and maintaining high trust in shared institutions. In order to maintain balance going forward and to integrate the (biased) portion of citizens online, binding corrective referenda can serve to enhance the ability of citizens to share political arguments and experiences that can counterbalance digitally-driven fragmentation.

As an answer to the main research question, which asked whether the digital environment de-legitimized the democratic municipal campaign and election process, we must therefore give a qualified reply. The digital environment both did and did not delegitimize the democratic municipal election process.

It did not delegitimize the election, because the municipal elections saw increasing party diversity, wins for local parties, and little influence of bots. For the Hague, an ideologically novel 'consensus' building party won big, even as there was further fragmentation in the seat allocation of other parties. Though it is possible to envision a link between increasing political fragmentation and ICT possibilities through an expanded and more personalized public sphere, this does not necessarily delegitimize democracy, and particularly not the current electoral campaign. ICT-driven fragmentation can moreover also be good for democracy, when properly balanced by shared trusted representative intermediaries. There will be some forthcoming evidence on which parties and topics generate more intolerance online than others, and their links to the small amount of bots witnessed, but we can already confidently say that such intolerance, if present, has not yet delegitimized the election process as a whole in the Netherlands. Intolerance and polarization might furthermore contribute to derailing a democracy, but in itself are never sufficient.

The digital environment did delegitimize democratic elections in that 'fake news' themes played a big role in elections themselves, and in the sense that the use of intransparant and potentially illegal micro-targeting increased, and is becoming easier by the day. There is a strong rise in intermediary companies and trackers, which means that there is an increasing power of corporate actors in public affairs. Public, high-trust media and political parties are meanwhile under threat from a lack of (digital) resources and an inability to adapt, while the conservative right-wing groups online proved particularly vulnerable to disinformation campaigns. We saw some echo chambers online, and though we saw quite some engagement online between different echo chambers, algorithmic selection biases against a diverse presentation of political viewpoints, and in favor of one particular direction. More insight is moreover needed based on data from big online infrastructural platforms, not just

to get more clarity, but also to ensure that citizens remain equipped to exercise basic autonomous democratic control over their shared environment.

Several problems in research thus require action. These stem primarily from the fact that much of the data that is needed to ascertain whether privacy/personal data is being used illegally for political purposes is behind closed doors, or owned by private companies. Supervisors, politicians, and researchers should have access to this information, if not the public at large.

Bibliography

Allcott, H., & Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, *31*(2), 211–236. https://doi.org/10.1257/jep.31.2.211

Arendt, H. (1967, February). Truth and Politics. *The New Yorker*, 54.

Bouwman, C. (2018, March 20). Ook hier willen partijen 'microtargeting.' *NRC*. Retrieved from https://www.nrc.nl/nieuws/2018/03/20/ook-hier-willen-partijen-microtargeting-a1596460

Commissariaat voor de Media. (2018). Nederlanders hebben weinig zorgen over nepnieuws. Retrieved September 28, 2018, from https://www.cvdm.nl/nieuws/nederlanders-hebben-weinig-zorgen-over-nepnieuws/#

de Cock Buning, M., Eljon, E., & Bune, J. (2018). *Reuters Institute Digital News Report Nederland 2018 Mediabedrijven en mediamarkten*. Retrieved from https://www.cvdm.nl/wp-content/uploads/2018/06/Mediamonitor-2018-Reuters-Institute-Digital-News-Report-Nederland-2018.pdf

Dillet, R. (2018, July 26). Facebook officially loses $123 billion in value. *TechCrunch*. Retrieved from https://techcrunch.com/2018/07/26/facebook-officially-loses-123-billion-in-value/

Economist Intelligence Unit. (2017). *EIU Democracy Index 2017*. Retrieved from https://infographics.economist.com/2018/DemocracyIndex/

Fletcher, R. (2018). *Misinformation and Disinformation Unpacked - Reuters Institute Digital News Report*. Oxford. Retrieved from http://www.digitalnewsreport.org/survey/2018/misinformation-and-disinformation-unpacked/

Foa, R. S., & Mounk, Y. (2017). The Danger of Deconsolidation: The Democratic Disconnect. *Journal of Democracy*, *28*(1), 5–16. Retrieved from https://www.journalofdemocracy.org/article/danger-deconsolidation-democratic-disconnect

Fredericks, B. (2018). Nearly half of Republicans say Donald Trump should have power to shut down media. *New York Post*.

Guess, A. (2018). *(Almost) Everything in Moderation: New Evidence on Americans' Online Media Diets* |. Retrieved from https://csdp.princeton.edu/publications/almost-everything-moderation-new-evidence-americans-online-media-diets

Guess, A., Nyhan, B., & Reifler, J. (2018). *Selective Exposure to Misinformation: Evidence from the consumption of fake news during the 2016 U.S. presidential campaign*. Retrieved from http://www.dartmouth.edu/~nyhan/fake-news-2016.pdf

Helbing, D., Zwitter, A., Frey, B. S., Gigerenzer, G., Hafen, E., Hagner, M., … Zicari, R. (2017, June 29). Will Democracy Survive Big Data and Artificial Intelligence? *Scientific American*. Retrieved from https://www.scientificamerican.com/article/will-democracy-survive-big-data-and-artificial-intelligence/

Jamieson, K. H. (2018). *Cyberwar : how Russian hackers and trolls helped elect a president : what we don't, can't, and do know*. Oxford: Oxford University Press. Retrieved from

https://global.oup.com/academic/product/cyberwar-9780190915810?cc=nl&lang=en&

Kim, T., Barasz, K., & John, L. K. (2018). Why Am I Seeing This Ad? The Effect of Ad Transparency on Ad Effectiveness. *Journal of Consumer Research*. https://doi.org/10.1093/jcr/ucy039

Lovink, G. (2008). *Zero comments : blogging and critical Internet culture*. Routledge. Retrieved from https://books.google.nl/books/about/Zero_Comments.html?id=WWjGQgAACAAJ&source= kp_book_description&redir_esc=y

Mounk, Y. (2018). *The people vs. democracy : why our freedom is in danger and how to save it*. Boston: Harvard University Press. Retrieved from https://books.google.nl/books/about/The_People_Vs_Democracy.html?id=H8tDtAEACAAJ &source=kp_book_description&redir_esc=y

Oudenampsen, M. (2018). *The conservative embrace of progressive values*. Tilburg University. Retrieved from https://pure.uvt.nl/ws/files/20023602/Oudenampsen_Conservative_12_01_2018_.pdf

Rawls, J. (2005). *Political Liberalism* (2nd Revise). Columbia University Press.

Snyder, T. (2018). *The road to unfreedom : Russia, Europe, America*. Crown/Archetype. Retrieved from https://books.google.nl/books/about/The_Road_to_Unfreedom.html?id=OeMxDwAAQBAJ &source=kp_cover&redir_esc=y

Urbinati, N. (2014). *Democracy disfigured : opinion, truth, and the people*. Harvard University Press. Retrieved from https://books.google.nl/books?id=3HbzAgAAQBAJ&dq=urbinati+democracy+disfigured&so urce=gbs_navlinks_s

Varieties of Democracy Institute. (2018). *Democracy for All?* Retrieved from https://www.v-dem.net/media/filer_public/3f/19/3f19efc9-e25f-4356-b159-b5c0ec894115/v-dem_democracy_report_2018.pdf

Verhagen, L., & Andersen, R. (2018, August 10). Klanten hebben recht op hun persoonlijke gegevens - maar krijgen ze die ook? We testten bekende bedrijven. *De Volkskrant*. Retrieved from https://www.volkskrant.nl/economie/klanten-hebben-recht-op-hun-persoonlijke-gegevens-maar-krijgen-ze-die-ook-we-testten-bekende-bedrijven~b7e6d7c8/

Vice News. (2016). We confronted Richard Spencer, the white nationalist known for using Nazi rallying cries. HBO. Retrieved from https://news.vice.com/en_us/article/wjzw3q/we-confronted-richard-spencer-the-white-nationalist-known-for-using-nazi-rallying-cries

Vidya Narayanan, Vlad Barash, John Kelly, Bence Kollanyi, Lisa-Maria Neudert, and P. N. H. (2018). *Polarization, Partisanship and Junk News Consumption over Social Media in the US*. Oxford. Retrieved from http://comprop.oii.ox.ac.uk/research/polarization-partisanship-and-junk-news/

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science (New York, N.Y.)*, *359*(6380), 1146–1151. https://doi.org/10.1126/science.aap9559

Zuiderveen Borgesius, F. J., Möller, J., Kruikemeier, S., Fathaigh, R. Ó., Irion, K., Dobber, T., … De Vreese, C. (2018). Online Political Microtargeting: Promises and Threats for Democracy, *14*(1). https://doi.org/10.18352/ulr