



UvA-DARE (Digital Academic Repository)

Data & Democracy

Political microtargeting: A threat to electoral integrity?

Dobber, T.

Publication date

2020

Document Version

Other version

License

Other

[Link to publication](#)

Citation for published version (APA):

Dobber, T. (2020). *Data & Democracy: Political microtargeting: A threat to electoral integrity?* [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 1

Introduction

Can you think of any strait, any Channel, that has the currents and variety of rough patches and changes of tide strong enough to match the upsets and the ebb and flow that accompany the working of elections? The whole situation is often changed by having to break off for a day or by night intervening and the merest breath of a rumor sometimes changes everyone's views. Often, too, for no apparent reason the turn of events takes you by surprise and at times even the people are amazed at a result as if it were not itself responsible. Nothing is more fickle than people in a crowd, nothing harder to discover than how men intend to vote, nothing trickier than the whole way in which elections work. (Cicero, Pro Murena, 62 B.C.)

Political actors have always tried to influence citizens, win elections, and gain power. Lawmakers have always tried to safeguard the integrity of elections and reign in the power of dominant political actors. Citizens were, and still are, expected to be exposed to a diverse collection of free-flowing political ideas, and, ideally, engage in autonomous deliberation. To bolster that ideal of autonomous deliberation, lawmakers introduced media and privacy laws, and regulations for political advertising.

In spite of regulatory efforts to protect citizens from influential political actors, technological developments have ushered in an era of *surveillance capitalism* (Zuboff, 2019): a digital and commercial system that collects, tracks, analyzes and leverages information about citizens to gain insights into how individuals differ and how these differences can be used to achieve some goal. While such information about citizens may seem meaningless when inspected on a small scale, patterns emerge when analyzed on a large scale. Knowledge about these patterns can give advantages to political actors seeking power.

As technology progresses, new methods of political communication have come up that may challenge the integrity of elections and the balance of power within a democracy. *Political microtargeting (PMT)* may be such a new method. PMT is a way to capture the attention of citizens who are on the one hand very *reachable*, because they carry their phones with them at all times, because their whereabouts are tracked, and because their personal data such as home addresses are collected on a large scale. But on the other hand these citizens are very difficult to get through to, because the media environment produces a cacophony of potentially attention-grabbing impulses. It is in this overstimulating landscape that knowledge about individuals and about what makes them tick translates into attention, and, most importantly, increases the chances of exercising influence.

Citizen behavior is being recorded and stored in commercial databases. Political actors, or their intermediaries, automatically sort through the databases to find information that sets people apart from the rest. If this information is deemed meaningful, it can be used to attempt to grab the attention of the citizen, and subsequently to persuade him or her into thinking or doing something in line with the political actor's interests. For

example, a sudden interest in baby materials signals young parenthood. This information can be leveraged into content that may make the citizen have a certain thought. A political party could tailor an ad catering to the hopes or fears of a young parent in the hope this would grab the citizen's attention, and make him or her think positively about that political party. This process is called political microtargeting (PMT). There is no 'fixed' way to microtarget: via direct mail, social media, face-to-face, et cetera. The appeal of microtargeting is that people are potentially more responsive to messages that are personally relevant than to generic messages. I will further conceptualize PMT later on in the introduction.

Scholars worry about PMT's threats to electoral integrity (Zuiderveen Borgesius et al., 2018, Gorton, 2016, Tufekci, 2014; Bennett, 2015; Barocas, 2012; Moore & Tambini, 2018), but knowledge about PMT is scarce, especially in a European context. This dissertation aims to assess if, how, and the extent to which PMT threatens electoral integrity. To solve that leading puzzle, we need to discover to what extent and under which conditions PMT is possible and prevalent in a European multiparty democracy, how voters perceive PMT, what effects are of PMT, and how PMT can be used to amplify manipulative information? This dissertation will address these key questions. More specifically, I first ask: *what barriers and facilitators for the adoption and use of PMT techniques do Dutch political parties perceive?* Then, I ask: *how does the Dutch electorate perceive PMT, and how do privacy concerns and attitude toward PMT influence each other over time?* In the next chapter, I ask: *what are effects of microtargeted issue-based messages on citizens' vote likelihood and issue salience?* Finally, I ask: *to what extent does a (microtargeted) deepfake meant to discredit a politician affect citizens' attitudes toward that politician and his party?* In the coming section, I discuss the emergence and conceptualization of PMT.

PMT Emerges

Commercial companies were the first to discover the potential benefits of tailoring ads to individuals. Showing consumers relevant ads should lead to higher *click through rates*, which, in turn, should lead to more purchases (Turow, 2012, Zuboff, 2019). Political parties in the US quickly, too, discovered the potential of data-driven tailored ads. They already sent personalized postal mail to potential voters, but data-driven (online) campaigning opened up a new range of possibilities to influence voters.

Success stories of political candidates that won an election (e.g. Obama 2008), often hailed the technological prowess of those campaigns. (e.g., Stirland, 2008; Murphy, 2012; Brennan, 2012). Competing campaigns, in following elections, would 'learn' from the success stories of their rivals and model their new campaign after their competitor's previously successful one (Kreiss, 2016). For example, for the US 2012 presidential race, Mitt Romney rolled out 'project Orca', a digital strategy modeled after (but meant to outperform) Obama's 'project Narwhal' (Kreiss, 2016). Not only

US campaigns started using more technology in their campaigns, Canada and the UK followed (Bennett, 2015; Anstead, 2017).

2016 was a pivotal year for the perception of digital tools in politics in general and PMT in specific. The (outcome of) the Brexit campaign and the election of Donald Trump as US president surprised many. Both the Brexit and the Trump campaign used PMT-techniques. News stories highlighted the role PMT played in both campaigns, suggested the technique was some panacea, uniquely able to persuade voters (Soares, 2016; Cadwalladr, 2017; Mahboob, 2019). While PMT's effectiveness was all but proven, campaigns in non-Anglo-Saxon, multiparty democracies also began to adopt PMT techniques (Drepper, 2017; Dobber, Trilling, Helberger & De Vreese, 2017; International IDEA, 2018).

In the Netherlands too, the political landscape started to turn increasingly innovative. The Dutch Green Party hired the services of US digital consultancy Blue State Digital (Hendrixx, 2016). The PvdA (social democrats) employed the services of one of the largest ad agencies in the Netherlands ('n=5'; Mebius, 2019). The largest Dutch party, VVD, hired Austrian data company *made2matter* for their campaign (Meeus, 2018). At the same time, Facebook connected more and more people, and, thus, collected more and more personal data of Dutch citizens. As did Google. Dutch campaigns could interact with citizens in ways unthinkable ten to fifteen years ago (Brants, 2006).

The Concept of Political Microtargeting

PMT evolves as technology evolves. In a relatively short period of time, PMT has been conceptualized in many different ways. As 'computational politics', for instance: "applying computational methods to large datasets derived from online and off-line data sources for conducting outreach, persuasion and mobilization in the service of electing, furthering or opposing a candidate, a policy or legislation" (Tufekci, 2014; p. 1). According to a description by Gorton (2016 p. 62), microtargeting "involves creating finely honed messages targeted at narrow categories of voters based on sophisticated combinatorial analysis of data garnered from individuals' demographic characteristics and consumer and lifestyle habits". Hersh (2015) does not explicitly define PMT, but makes it clear that it involves "targeting voters based on attributes captured in campaign databases" (p. 205). Rubinstein (2014; p. 882) describes microtargeting as a "form of political direct marketing in which political actors target personalized messages to individual voters by applying predictive modeling techniques to massive troves of voter data." PMT may bring about associations of the offline 'ward politics' decades ago, when local political campaigns were very knowledgeable about the people living in their districts. The internet brought what Kreiss (2016) refers to as 'networked ward politics', where campaigning is different from the old days in three ways: "first, in terms of campaigns leveraging 'people as media' in the course of field canvassing and social media appeals; second, in terms of campaigns meeting citizen expectations to determine their own levels of often highly temporal

political engagement; third, in terms of campaigns increasingly appealing to citizens as individuals and not as representatives of demographic groups” (p. 217).

Green and Gerber (2008; p. 56) described ‘microtargeting databases’ that enabled political campaigns “to tailor mailings to match the ostensible political viewpoints of different segments of the electorate.” This was in 2008, and understandably focused on offline microtargeting. There is also research that strictly focuses on online political microtargeting. Zuiderveen Borgesius and colleagues (2018, p. 82) described online political microtargeting as “a type of personalized communication that involves collecting information about people, and using that information to show them targeted political advertisements.” I would argue that the offline versus online distinction is no longer useful. Because inferences and predictions are based on multiple data points that are not necessarily ‘online’ or ‘offline’, such as location data, or consumer data, the ‘real world’ and ‘cyberspace’ have become intertwined and the distinction has become confusing. Even if the microtargeted message is delivered online or offline, we should simply speak of microtargeting. Especially in light of developments around tailored advertising on smart tv’s (Perakakis & Ghinea, 2015) and synced advertising (Segijn, 2019), which, in practice, imply that a person can be targeted with the same ad across devices (phone, tv, billboard) and across locations (at home, on the road, at work).

I do, however, use Zuiderveen Borgesius’ et al (2018) definition of PMT in this dissertation, but lose the distinction about online or offline microtargeting. Other definitions are either too specific about in which kind of databases personal information is being stored (e.g. Hersh, 2015), too ‘exaggerating’ of the “sophisticated” techniques used (e.g. Gorton, 2016), too limiting about the amount of data that can be used (Rubinstein, 2014), or too limiting in the purposes for which PMT can be used (Tufekci, 2014).

In a similar vein, there are several different terms for ‘political microtargeting’ that overlap to a great extent. For example, ‘Political behavioral targeting’ entails tailoring political messages on the basis of past behavior of citizens. ‘Psychometric’ or ‘Psychographic’ targeting means tailoring political messages to personality traits of the receiver. ‘Issue-based targeting’ is about tailoring a message to the issues that are perceived as most salient to the target audience. All terms refer to some a-priori data collection, and subsequent tailoring to a certain type of characteristic, but the specific type of characteristic differs per term. I would argue that political microtargeting is the better term, as it functions as an umbrella that covers all other terms as well. Moreover, political microtargeting (PMT) also captures the fusion of two or more terms. For instance, when a political advertiser targets people on their most salient issue, *and* then further tailors the message to the personality traits of the receiver (e.g. an introvert that cares most about education receives different messages than an extravert that cares most about education).

Actors

PMT is a dynamic process, involving different interdependent actors. The main actors are social platforms, political advertisers, data brokers, and citizens. The lion's share of this dissertation focuses on political advertisers and citizens. Social platforms are important too, but to a lesser degree. Data brokers are not the focus of this dissertation.

Except for the data brokers, the actors take on different roles at the same time. *Social platforms*, such as Facebook, are first of all merchants. They sell insights and ad space. Second of all, they are intermediaries: connecting advertisers and citizens. Third of all, they are captivators: engaging their users and giving them a good experience. It is no wonder that Facebook charges more money to Republican advertisers who want to target Democratic voters than if they wanted to target Republican voters (Ali et al., 2019). These incongruent messages put social platforms like Facebook in a difficult position. On the short term, Facebook wants to connect advertisers and users. But on the long-term Facebook wants to keep on engaging its users so Facebook can continue to act as a merchant: selling insights on and access to the attention of active users. Incongruent messages potentially give users a worse experience than congruent messages (see congruence theory (Aaker, 1999)) and, thus, something should balance or hinder this (e.g. extra costs for the advertiser). In essence, these three roles have turned social platforms into political actors.

Political advertisers can collect and analyze data themselves, or hire data brokers or social platforms to do that for them. Due to the easy-to-use social platforms, any individual can be a political advertiser. As such, this is a diverse group of legitimate domestic advertisers, such as political parties, individual candidates, NGO's, or concerned citizens, but the group also consists of domestic or foreign actors who may have less legitimacy.

Data brokers continuously track the citizen, collect data about their behavior, buy and combine additional datasets and clean the data. On the basis of this information, data brokers infer, predict and profile citizens. Individual profiles are continuously updated. Data brokers operate in the background. They provide the fundamentals upon which microtargeting efforts rest.

Citizens. All other actors want something from the citizen. Be it attention or engagement (social platforms), votes (political advertisers), or information (data brokers). The actors hope to gain money, and/or power from citizens. Citizens do not live in a vacuum, and microtargeted messages are only part of the political information they consume. Citizens play an active role in submitting information that is later used to influence them or people that look like them: essentially feeding the surveillance capitalism industry (see Zuboff, 2019).

PMT versus traditional ways of audience targeting

PMT, at first glance, may look similar to more traditional ways of audience targeting. But PMT is fundamentally different. Regular targeting (i.e. not *microtargeting*) is based on group level data, PMT on individual level data. Say a political party wants to reach likely voters. A traditional approach could be to buy an advertisement in a magazine read by a specific group of people (e.g., right-leaning, young urban professionals). While the ad would likely be seen by many potential voters, the ad would also reach many citizens who are never going to vote for the advertising party, or people who are not going to cast their vote at all. Apart from being considered inefficient, these traditional ways of advertising can also be considered less effective. The potential voters reached by the magazine ad do not form a homogeneous group. The group members likely care about a plethora of different political issues. Some care about immigration, some about crime, others about taxes, and so on. The problem is that if a political advertiser reaches this heterogeneous group of potential voters, they can never address all the different things important to all the different individuals. In fact, the advertiser needs to have (a lot of) detailed information about individuals' characteristics and preferences. It is not enough to know or predict that certain people are likely voters. Rather, the advertiser should know what makes people tick, and this is impossible to learn on the basis of group information. A different *non-microtargeting* approach of many political campaigns is to look at election results in specific neighborhoods and infer which neighborhoods or which streets have potential and which do not. While this approach is arguably more efficient than the magazine ad, it still does not acknowledge the heterogeneity of the high potential neighborhood or street. All residents (or rather, a large chunk of the group) may consider voting for one specific party. But they have their own different reasons, emotional and rational, for that.

Instead of using group level data, PMT uses individual level data as *input*. This way, PMT leverages individual characteristics with ads specifically tailored to those characteristics. The idea is that this would make the receiver especially susceptible to the ad. This does not mean that a political campaign needs to craft millions of individually tailored messages. A political campaign could identify specific people who share the characteristics that would make them susceptible to a specific message, group them, and microtarget them. See for example the ads of the Dutch Socialist Party below. The first ad calls for higher salaries, the second ad for lower value added tax rates. One could imagine that the first ad was aimed at employed people, and the second ad might be aimed at people who are not employed (e.g., pensioners, or people collecting unemployment benefits). This is by no means a granular form of microtargeting, but it illustrates how messages can be tailored to influence different people.



Image 1 - Ad calling for higher salaries



Image 2 – Ad calling for lower VAT

PMT is fueled by inferences from and predictions of human behavior. While ‘microtargeting’ suggests a certain precision, a political advertiser targets on the basis of perceptions. When a specific citizen receives a tailored advertisement from the Socialist Party, this is because the data make the advertiser perceive that citizen as a likely Socialist Party voter, not because she actually is one (see Hersh, 2015). The better the data and the analysis, the better the match between advertisement and citizen. But the more specific an ad is tailored, i.e. the more inferences or predictions that form the basis of the perception, the more likely an advertiser ‘gets it wrong’. Political advertisers get penalized for ‘getting it wrong’ (mistargeting), which leads to *less* votes (Hersh & Schaffner, 2013).

Threats of PMT to electoral integrity

While technology has enabled political actors to directly communicate with citizens, it has also created a certain distance that follows from the opacity that surrounds PMT. A political actor can run a microtargeting campaign without ever leaving her office. A voter can come into contact with tailored political messages without ever seeing a political actor. This distance is fertile ground for unfair practices that could in the short or long run threaten the integrity of democratic elections.

In the 1900's, political actors made house visits on a relatively large scale to gauge vote intentions and to influence potential voters (De Jong, 2005). Back then, unfair practices, such as disinforming supporters of competing parties about issue positions, or the election date would have been unlikely to occur on a large scale because of the personal aspect of house visits: lying is more difficult when you do it to someone's face. The introduction of the television ushered in a new form of political campaigning. And while television is not personal, it is transparent: anyone with access to a television can scrutinize the same political message, which makes it difficult to conduct unfair practices people without anyone finding out. Technology brought back the opacity of campaigns that had disappeared with the rise of tv campaigns, and technology allowed political actors to *personalize* messages without the need for personal contact. Moreover, technology enables political actors to operate on a large scale. It is against this backdrop that especially the negative potential of PMT has received scholarly attention: because technology makes it so easy to lie, cheat and otherwise threaten the integrity of democratic elections.

Scholars have identified several ways in which PMT can threaten the integrity of elections. These potential threats are generally identified by normative analysis and not yet supported by empirical research. Before we can determine whether a threat is more than just *potential*, we need empirical research. I will discuss the major potential threats to electoral integrity (privacy, intellectual privacy, manipulation, fragmentation of the public agenda and changing political power structures) and indicate how the empirical research in this dissertation contributes to a better understanding of these potential threats. These major threats are derived from the Council of Europe's Venice Committee (2002), that identified five principles underlying good practices in elections: there should be universal, equal, free, secret and direct suffrage.

Privacy

Privacy relates to the Venice Committee's principles of free and secret suffrage. Personal data fuel PMT. Learning more about people's political preferences, their susceptibilities, and how to reach them can be at odds with citizens' privacy. Privacy infringement opens the door to voter manipulation (see below). And where personal data are being amassed there is always a risk of data breaches. The General Data Protection Regulation (GDPR) is meant to make sure that the collection and processing of personal data in general occurs in a fair and transparent manner, and that in particular data on political opinions is collected and processed under strict conditions. But the GDPR does not explicitly mention PMT, and there are issues regarding the protection of political speech that make the implications of the GDPR regarding PMT unclear (see Dobber, Fathaigh & Zuiderveen Borgesius, 2019). In the second chapter of this dissertation, I explore how political parties struggle with the obstacles to data collection that follow from the GDPR. In chapter three, I model the Dutch citizens' attitudes toward PMT and their privacy concerns over time. In chapter four, I report a field experiment

where I collect, and use personal political data (after receiving informed consent) to microtarget citizens. In chapter five, I leverage information about people's religiosity to tailor a deepfake to them personally. Together, the chapters show the role of privacy in political campaigning and how a lack of privacy can be (mis)used to influence citizens.

Intellectual privacy

Intellectual privacy relates to the principles of free and secret suffrage. Scholars also worry about the disappearance of one's room to privately and personally deliberate (Richards, 2015; Reiman, 1995). A citizen needs 'intellectual privacy' to take an informed voting decision (Richards, 2015). Constant surveillance severely threatens citizens' room to think freely (see Reiman, 1995), which limits citizens' capacity to deliberate (Dawes, 2014). And constant surveillance is precisely what is needed to identify individual differences and susceptibilities to which political actors tailor their messages. This threat to intellectual privacy could lead to chilling effects (citizens refrain from certain behavior because they perceive they are being monitored [Marder et al., 2016; Penney, 2016]). But the idea of intellectual privacy rests to a great degree upon the perceptions of citizens themselves. If citizens do not feel concerned about their privacy, it is very unlikely that they would suffer any negative consequences from this surveillance. Therefore, in chapter three, I examine the (intellectual) privacy perceptions of Dutch citizens and indeed find a downward spiral that over time can result in undesirable behavior of the voter from a democratic point of view (e.g. chilling effects). Chapter three focuses on the conditions under which certain undesirable behavior can occur due to negative privacy perceptions, but not on the actual behavior itself.

Voter manipulation

Voter manipulation relates to the principle of free suffrage. Where a newspaper advertisement is visible for all, microtargeted ads are visible for a few. Social platforms have made an important step with the introduction of ad libraries that show political advertisements running on the platform. But these libraries still have important pitfalls (see Leerssen et al., 2019; Tromble, Jacobs & Louwerson, 2019). Political advertisers can easily set up accounts that cannot be traced back to them and use those accounts to spread disinformation or misleading content without risking journalistic scrutiny. Moreover, ad libraries provide only limited information about who was reached by an ad, but no detailed information on targeting criteria used for an advertisement. Besides the opacity that surrounds PMT, there is an information asymmetry between the political advertiser and the citizen (Tufekci, 2014). Oftentimes, citizens are unaware that the ad they see is tailored to them personally. These circumstances open the door for deceit and manipulation. An advertiser could misleadingly present itself as a one-issue party, championing the issue most important to the receiver (Zuiderveen Borgesius et al., 2018). An advertiser could also blatantly lie or leverage the vulnerabilities of the citizen (Susser, Roessler, & Nissenbaum, 2019).

In chapter two of this dissertation, I explore how Dutch political parties innovate and adopt a context-dependent form of PMT. I learned, for instance, that political parties used *dark posts*, which can only be seen by the targeted group and no-one else: fertile breeding ground for deception. Chapter two also addresses the ethical and legal dimension of the use of PMT-techniques as perceived by Dutch campaign leaders. In chapter four, I measure the effectiveness of PMT ads, which is an important pillar of the PMT's manipulative potential. After all, if PMT ads are not effective, it is difficult to see how they could do harm. In the fifth chapter I show how PMT can be used to amplify the effects of a deepfake on political attitudes. The idea is: not every person is susceptible to the same disinformation. PMT-techniques can help malicious political actors reach the 'right' citizens with the 'right' piece of disinformation, and manipulate citizens' attitudes and opinions.

Fragmentation of the public agenda

Fragmentation relates to the principle of equal suffrage. Agenda setting theory describes how news media, by the topics they cover, influence what societal issues citizens find important (McCombs and Shaw, 1972). When people get see microtargeted personalized advertisements, this agenda becomes less 'public' and more 'personal'. As a result, public deliberation becomes increasingly difficult because people assign weight to a more diverse range of issues. Contributing to the fragmentation of the public agenda has been perceived as a threat in the Netherlands long before the advent of PMT. It was Dutch statesman Johan Thorbecke who, in the 1800's, warned his fellow politicians against catering to local issues, "*esprit de clocher*", and worked to safeguard the "general interest instead of particular interests" (Kaal, 2016, p. 489-490). Thorbecke went as far as drawing deliberately heterogeneous districts (from 1848 to 1918 the Netherlands had a district voting system), to protect the public agenda against local and sectional issues (Kaal, 2016). Different from 1848, people now are likely to get their political information from diverse, public, sources (e.g. public broadcasters: Trilling & Schoenbach, 2013; Trilling & Schoenbach, 2015). But the fragmentation of the public agenda becomes more salient when information diets become increasingly personalized (for instance in combination with personalized news recommendations), or with synced advertising (see Segijn, 2019).

In chapter two, I find that campaign leaders did not perceive the fragmentation of the public agenda as a threat that follows from PMT. Instead, they viewed PMT as an opportunity to convince or educate low-knowledge citizens. In chapter four, I examined whether receiving a political message tailored to the individual participant's most salient political issue would increase the personal salience of that issue, and decrease the salience of the other, less-salient, issues.

Political power structures

Political power structures relate to the principles of equal, free, and secret suffrage. PMT can affect political power structures in society. Large platforms, in essence, are political actors themselves. Platforms can influence politics by using their own infrastructure for political gain. The ‘61-million-experiment’, conducted by Bond et al. (2012) in cooperation with Facebook, showed that a small adjustment of someone’s timeline (presenting a social message reminding people to vote) could increase turnout significantly. The effect was small (.4%), but this translated into 360.000 extra votes. It is easy to see how powerful actors could help friendly or hinder hostile politicians getting elected.

There is a risk that elected officials will refrain from regulating the platforms or the techniques they themselves depend on for their re-election. A conflict of interest between elected officials’ legislative tasks on one hand, and their political tasks on the other hand, was signalled by Hersh (2015) in relation to the use and availability of data in the US. Hersh found politicians who appeared to propose measures to collect voter data, because it could help their campaigns in the next election. Hersh also found politicians actively rejecting privacy protection measures because their political campaigns need such data. This is not a strictly US phenomenon. Another illustration can be found in Australia. There, politicians have exempted themselves from the Australian Privacy Act. Citizens’ privacy should be protected, but not at the cost of the lawmakers’ election campaigns (Australian Privacy Act, 1988), it seems. In Spain, politicians tried to drastically lighten the General Data Protection Regulation’s relatively stringent requirements on data collection and processing for their own political purposes, but failed¹ (Gil Gonzalez, 2019).

Domestic actors as well as foreign actors (e.g., adversary nations) can use microtargeting techniques to more efficiently spread disinformation and meddle in democratic elections. Digital disinformation capabilities could, potentially, affect electoral outcomes in line with the political goal of the disinforming actor. In reality though, so far, it is doubtful and at best difficult to determine whether disinformation campaigns had any impact on electoral outcomes in large democracies such as the US (Guess et al., 2018; Bail et al., 2019; but see Jamieson, 2018), or smaller democracies such as Madagascar (Schwartz & Borgia, 2019). But since the US and Madagascan elections, technology has progressed and new disinformation modes, such as deepfakes, have surfaced. Deepfakes are a remarkable new mode of disinformation as they can make it seem as if a politician says or does something, while in reality the politician never did so. Deepfakes are potentially more powerful than bots, trolls, and false news stories because deepfakes are videos that deceive the eyes and the ears. Deepfakes can give malicious actors power over elections because they can sow confusion, doubt and distrust about what is real

1 The Spanish Constitutional Court nullified the Spanish law.

and what is not (Vaccari & Chadwick, 2019; Karpf, 2019). Moreover, deepfakes can empower malicious actors because they can potentially influence citizens to behave in ways aligned with their political goals. To make things worse, PMT techniques could potentially amplify the effect of a deepfake, by enabling malicious political actors to tailor a deepfake to the receiver.

In chapter five, I examine the effect of a political deepfake on citizens' political attitudes. Moreover, I explore whether microtargeting techniques can amplify its effects. In chapter two, I study the perceptions of campaign leaders toward the use and adoption of PMT techniques. I do not find a dependence on Facebook, but almost all parties use the platform and have been using the platform more intensely over the years. There is a dynamic between campaigns where they look at the actions of competitors and do not want to be at a disadvantage. If a competitor buys a lot of ads through Facebook, other parties will at least consider doing the same. Especially when that competitor seems to be successful due to their online strategy (see also Kreiss, 2016).

Dissertation context

PMT research overwhelmingly focuses on the US (Nielsen, 2012; Kreiss, 2012; Kreiss, 2016; Hersh, 2015; Bimber, 2014; Gorton, 2016; Endres, 2016; Endres, 2019; Haenschen & Jennings, 2019). But from a global perspective, the US is an outlier. The incomparable budgets, the lax regulation, and the electoral system make that US-based findings cannot necessarily be applied to other contexts. In this light, the Netherlands is a more representative case: a multiparty democracy², subject to a relatively stringent data protection regime regarding information about political preferences, and a liberal regulation regime regarding political advertising.

The Netherlands has an open list, one district, proportional representation (PR) electoral system. Elections are competitive, but, different from first-past-the-post systems, the Dutch parties may ideologically compete with their opposites, but in election campaigns parties also compete with their ideologically close peers (e.g. Dutch leftist parties compete for the same voter as their fellow leftist parties, see Rekker & Rosema, 2019). How useful is PMT in this context? Can PMT help advertisers discern CDA voters from ChristenUnie and SGP voters, for example?

Dutch political parties have small budgets (i.e., less than 10 million euro). How does that impact the use of PMT? Can low-budget campaigns even use PMT techniques or is this something that only the larger parties can do? What does this mean for the level playing field? In the Netherlands, election turnout is relatively high (almost 82% in the last national parliamentary election): would Dutch advertisers focus only on persuasion, or use PMT-techniques to further increase turnout? These questions are all

2 In 2020, there are 13 elected parties in parliament and 2 one-man parties that have splitted off.

studied in this dissertation. Together, these factors mean that PMT takes on a different shape in the Netherlands than it does in the US, and that Dutch-based findings are better suited to be translated to other European multi-party contexts than US-based findings are.

Studies

The first empirical chapter, chapter 2, aims to answer the following key question: *what barriers and facilitators for the adoption and use of PMT techniques do Dutch political parties perceive?* Chapter 2 relies on elite-interviews with campaign leaders of Dutch political parties, collected in the run-up to the Dutch national election of 15 March 2017. Answering this key question requires a qualitative approach for two main reasons. First, because there are only a handful of people in the Netherlands best suited to talk about how political campaigns perceive PMT and the conditions under which the technique is being adopted and used. Second, because this question is about understanding the perception of and mechanisms behind PMT. These are explorative, in-depth questions that cannot be sufficiently understood via quantitative methods (see also Boeije, 2005).

I then focus on how citizens perceive PMT over time, by answering the following question: *how does the Dutch electorate perceive PMT, and how do privacy concerns and attitude toward PMT influence each other over time?* This second study, chapter 3, relies on a three-wave panel study ($N = 879$) collected by CentERdata and the University of Amsterdam. A three-wave panel study is useful to track attitudinal and perceptual changes over time and enabled me to model the reciprocal dynamic of Dutch citizens' attitudes toward PMT and their privacy concerns over time.

After that, I examine the effects of PMT, by asking: *what are effects of microtargeted issue-based messages on citizens' vote likelihood and issue salience?* This third study, chapter 4, is a field experiment, carried out before and after the municipality elections of March 21, 2018. The study relies on a round of data collection before the election ($N = 124$) and a repeated measurement right after the election ($N = 86$). A field-experimental set-up was needed to keep the measurement of PMT's effects as natural as possible. To be able to isolate the effect of the stimulus, participants were subjected to a pre-election and post-election measurement.

Finally, I explore how PMT be used to amplify deepfake disinformation, by asking: *to what extent does a (microtargeted) deepfake meant to discredit a politician affect citizens' attitudes toward that politician and his party?* This fourth study, chapter 5, is an online experiment and relies on one round of data collection ($N = 278$). A deepfake (video and audio) was created from scratch for this study. An online experiment was the best approach to examine the amplifying role of PMT in terms of both feasibility (being able to target people with specific characteristics) and ethics (an online experiment is

a controlled environment, nullifying the risk of the deepfake spreading; a more natural setting would bring the risk of actually disinforming the public). All studies were conducted in the Netherlands.

In this dissertation, there may be some overlap between chapters' introductions and theoretical frameworks, as the chapters were written as standalone articles.