HOW TO FOLD A FLAG

Untangling the #nexit movement on Twitter

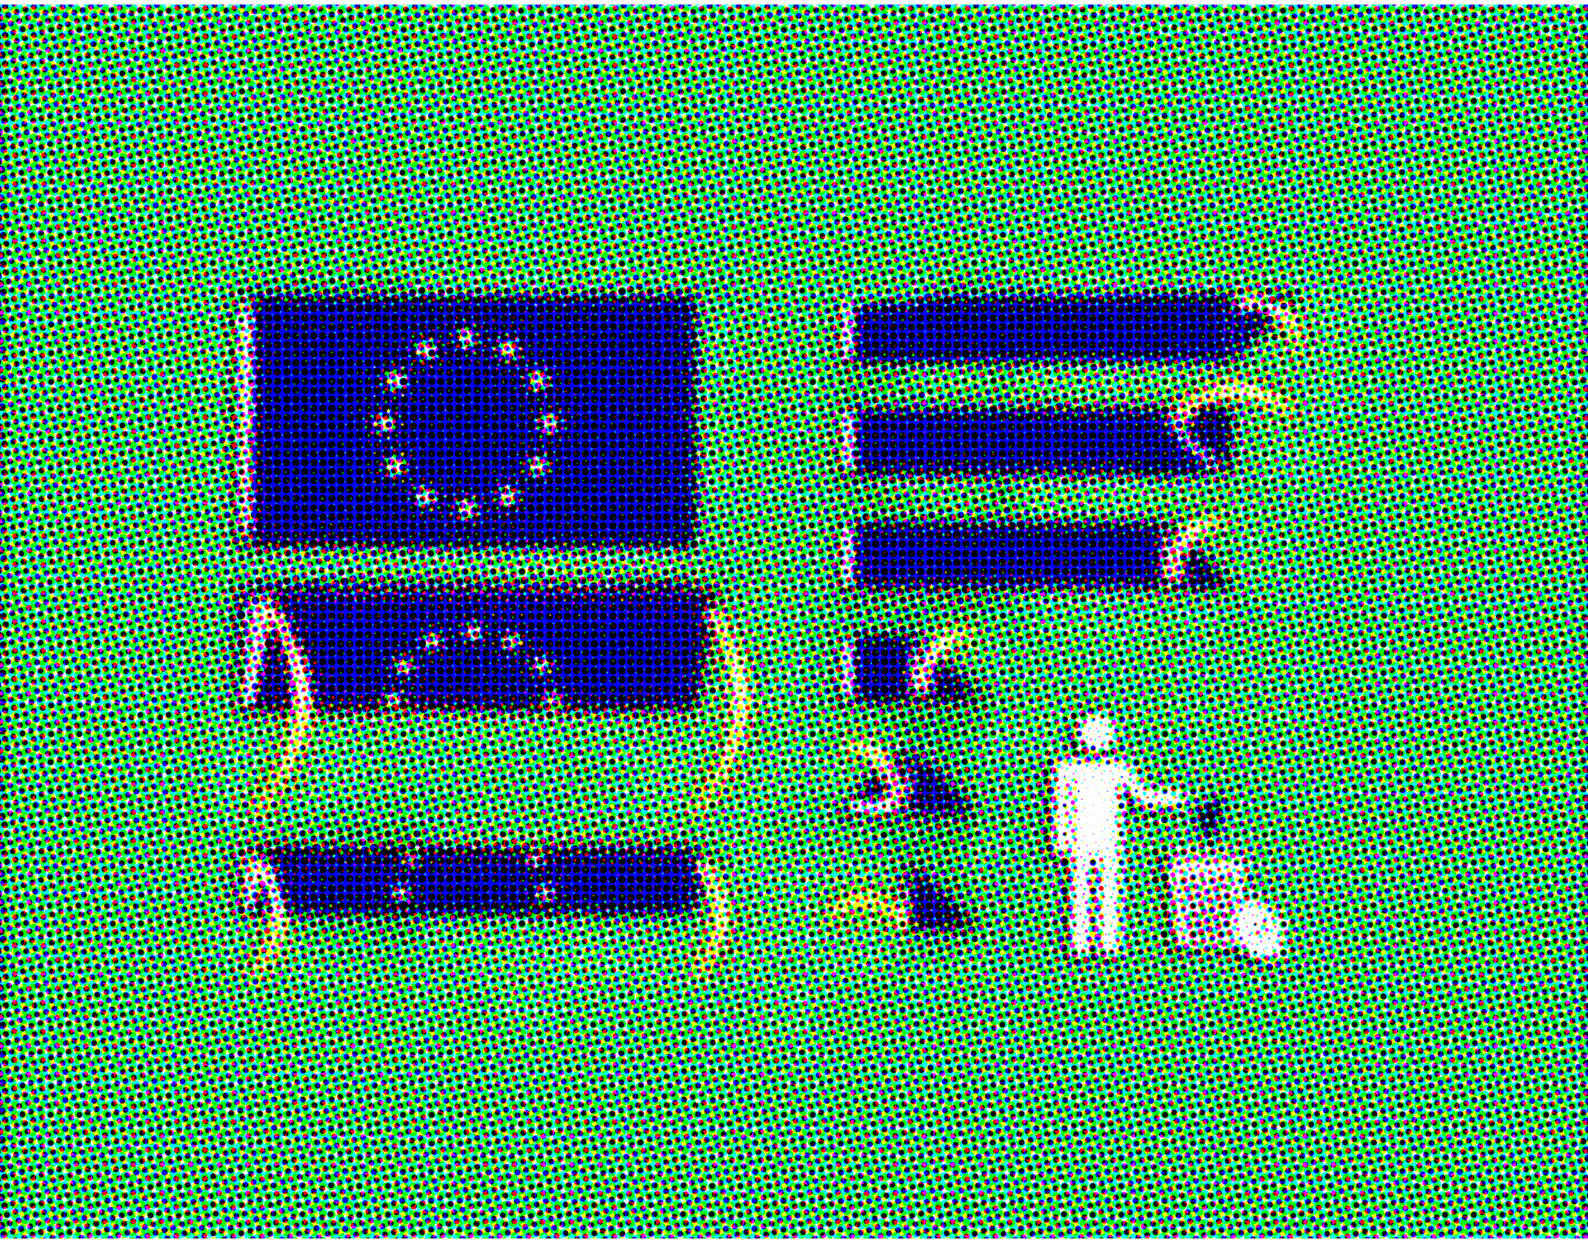

Utrecht University
MA thesis New Media & Digital Culture
Student name: Jeroen Bakker (6914721)
Supervisor: dr. Mirko Tobias Schäfer
9.730 words
23 October 2020

Table of contents

Abstract

Twitter has become a prominent arena for online political debate, and has provided a platform to a wide range of political voices. Political dissidents utilize the platform to spread news from "alternative media" and follow like-minded opinion leaders. Though studies into such political topic communities are plentiful, they often fail to acknowledge the role of socio-technical platform dynamics in the formation of such ideologies: the many different aspects of surrounding the platform that impact how information is created and circulated. This research seeks to address this gap by researching one such topic community explicitly through the lens of platform dynamics. By analyzing a large Twitter dataset of the cluster surrounding the hashtag #nexit– a vocal group of Dutch anti-EU Twitter accounts – I show that reactionary discourse thrives on the platform as it fits its inherent purpose: to drive user engagement. Through a variety of explorative data analysis methods, I demonstrate that the popularity of #nexit is strongly tied to news events and actualities, and that populist politicians exploit the anti-authoritarian connotations of the hashtag to engage the most radical members of their vote base. I argue for further acknowledgement of the role of platform dynamics in shaping extreme ideologies, not only in academic research but also in policymaking and platform design.

Introduction

Over the past decade, Twitter has established itself as the dominant social media platform for political discussion, granting politicians, journalists, experts and citizens the opportunity to debate on the same platform (Stieglitz & Dang-Xuan 2013; Enli & Skogerbø 2013). Simultaneously, the micro-blogging platform has drawn criticism for facilitating political extremism and polarizing communities (e.g. Alizadeh *et al.* 2019; O'Neil 2017). This phenomenon has led to a large number of studies into the sociological and political dynamics within these debates (e.g. Williams, McMurray & Lamber 2015; Feller *et al.* 2011; Conover *et al.* 2011). Critics have pointed out that this seemingly democratic online debate is strongly influenced by how the platform holders design and govern social media (e.g. Gillespie 2018; Bucher & Helmond 2016). Social media platforms could be considered black boxes to a certain extent; the algorithms that influence daily communication on social media are well-kept company secrets that are essential to the companies' business models. The effect of (largely hidden) social media platform dynamics on the formation of ideological communities has been a blind spot for academics who study such groups, with many choosing to focus on the visible aspect of platforms instead, such as functionality (e.g. Bossetta 2018), moderation (e.g. Gillespie 2018) content filters (e.g. Humphreys 2015; Stier *et al.* 2017) and user agreements (e.g. Van Dijck, Poell and De Waal 2018). Studying the opaque web of dynamics that social media engage in is difficult and requires accepting a certain level of uncertainty, as only the in- and outputs of a platform are known (Bucher 2018). Explorative research can however contribute to building a framework for further analysis.

In this thesis, I seek to make explicit these rather inconspicuous, opaque or hidden dynamics that emerge from Twitter's platform structure and have a profound impact on how communities form and information spreads on Twitter and beyond (Bucher 2018). The implications of these platform dynamics are critically studied through an extensive 'micro-study' of one ideological cluster on Twitter: the vocal community surrounding the Dutch anti-EU hashtag #nexit. In the analysis of an extensive dataset, we discover an exceptionally tight-knight, homophilic cluster of accounts that could be considered an echo chamber (Van Geenen *et al.* 2016). We can already point to some well-known elements of Twitter that might have contributed to this: Twitter's algorithmic discovery engine picks up on frequently mentioned accounts, terms and hashtags, and connects them based on similarities and co-occurrences; opposing voices can be muted, blocked or spun to fit the anti-EU narrative. The consequences – the radicalization of a group of people who rarely interact with users from outside of their echo chamber (Mansour-Ille 2020; Grover & Mark 2019; Dauber & Ilter 2019), the attention the "movement" is given by mainstream media, the legitimization of these ideas through this mainstream attention – are very real. Putting the blame solely on Twitter's platform design, however, is an obsolete, techno-deterministic approach that does not consider the multitude of other dynamisms at play in constructing an ideological hub. A better understanding of what role these platform dynamics play in facilitating a breeding ground for extreme beliefs is therefore essential.

Sentiments expressed within the #nexit community point towards neo-nationalism, a reactionary ideology postulating nationalist and nativist positions that rose to prominence in the 2010's (Eger & Valdez 2014; Antonsich 2017; Bergmann 2020). Social media are considered to play a role in lending this fringe ideology a platform in recent years (Postel-Vinay 2017; Fuchs 2020), though explicit connections between the two have yet to gain widespread academic attention. I am not so much concerned with the specific ideology of the #nexit cluster itself, but rather how its extremism is facilitated and amplified through Twitter's dynamics. #nexit, in this, serves as a case study of a topic community that has evolved into an ideological echo chamber. Drawing from cultural analytics (Manovich 2009) and digital methods (Rogers 2013), I make visible patterns that reveal the dynamics within this community. The central research question is: How do platform dynamics affect and contribute to shaping the #nexit topic community on Twitter? The insights of this study contribute to an academic framework for further research of such online political extremism by making explicit the hidden dynamisms that arise from a rich Twitter dataset.

The opaque web: platform dynamics

This study approaches the Twitter community surrounding #nexit through the lens of social media platform dynamics. In short, platform dynamics are all different dynamisms that arise from a platform structure. What platforms afford users to do, and what they don't, dictates how we socialize, consume, discover and communicate (Bucher & Helmond 2016; boyd 2010). Platform structures have infiltrated many layers of modern life, so much so that we could state that we live in a "platform society" (Van Dijck, Poell & De Waal 2018).

Within this society, the large corporations that govern the platforms have far-reaching influence on daily life, as the engineers at these companies are responsible for the algorithms that process the data, information and content and dictate what users can see and do (Gillespie 2018).

Social media platform providers are most often concerned with keeping users engaged, as engagement generates more user data that can be sold to third parties or used to spot advertisement opportunities (Wieneke & Lehrer 2016; Estrada-Jiménez *et al.* 2019; Ruckenstein & Granroth 2019). These practices have led to much discussion of the ethics of social media companies, with widespread attention being brought to issues such as the infringement of user privacy (e.g. Smith *et al.* 2012; Beigi 2018), addiction (e.g. Blackwell *et al.* 2017) and political interference (e.g. Woolley 2016; Badawy *et al.* 2019). What all these studies have in common, is that they focus on the in- and outputs of social media platforms. The processes in between are notoriously difficult to study, as platform designers hide the intricate inner workings of the medium behind user-friendly interfaces that serve to streamline the user experience and help to make sense of enormous amounts of data, but simultaneously obscure what happens behind the scenes (Van Dijck, Poell & De Waal 2018). These interfaces are designed to become ubiquitous; they disappear into the background and thereby create the illusion that the content is 'directly' presented to the user. This is reflected in many studies of political communication on social media, which often only focus on the content of the communication and fail to sufficiently acknowledge the many dynamics that influence how that content is distributed, filtered and interacted with (e.g. Stier *et al.* 2018; Enli & Skogerbø 2013).

Then how do we bring these invisible dynamics to the light? Eriksson *et al.* (2016) offer a solution with their explorative research into Spotify. Through a series of inquisitive experiments, they try to lay bare the internal processes of the music streaming platform. Their methods are often based on educated guesses and gut feelings; their material are the in- and outputs of the platform. Instead of merely studying those as separate entities, they try to pick apart what happened *between* the in- and output for the results to turn out this way. While Spotify primarily consists of interactions between the user and the content on the platform, Twitter is a social network platform that facilitates interactions between users, meaning that socio-cultural practices are brought into the online sphere. The "in between" of the data in- and outputs is therefore strongly influenced by social conventions and phenomena. In the following, I will reflect on those aspects of political communities on Twitter.

Echo chambers and opinion leaders

The #nexit community is an example of an ideological hub (Wieringa *et al.* 2016), although we can also speak of an *echo chamber* based on the high level of interconnectedness between accounts. Such communities are characterized by a high degree of homophily; almost all accounts within the cluster share the same views and actively seek out information that confirms their beliefs (Becker 2015: 143-144). "Echo chamber" is often used interchangeably with the term "filter bubble", but these concepts have different denotations. A filter bubble supposedly categorizes users through the platform's algorithms and places them in an isolated social environment, where they only ever come into contact with content that confirms their pre-existing beliefs. This concept was coined in a TED Talk by Pariser (2011) and has seen some academic use (e.g. Flaxman, Goel & Rao 2016; Spohr 2017), but was quickly criticized for not accurately reflecting how social media function in practice (e.g. Bruns 2019; Wieringa, De Winkel and Lewis 2017; Guess *et al.* 2018; Zuiderveen Borgesius *et al.* 2016). Wieringa, De Winkel and Lewis challenge the concept of the filter bubble by demonstrating that there is often a lot of overlap between online topic communities with differing political views. Accounts from different communities often share many of the same articles and sources, and react to the same popular accounts. Furthermore, platform elements such as the "trending topics" section on Twitter and "Discover" sections on other platforms actively encourage users to seek out content and accounts from outside of their usual circle. The notion of "echo chambers" acknowledges that it is inevitable for opposing information to find its way into a community.

Echo chambers are very much the result of socio-technical platform dynamics. The affordances that Twitter grants its users, such as blocking, liking and muting, give them a chance to shape their own social environment on the platform, which is amplified further by algorithms that attempt to recommend similar content (Alini *et al.* 2017; Bucher 2018). Engagement can also be driven by conflict, however. Negative sentiments have been shown to lead to more interactions (Park 2015; Vargo & Hopp 2020; McCreery & Krach 2018). Information that challenges a community's beliefs, such as news articles or content from other accounts, often leads to a vocal reaction and thus more content and user data that gets fed back into the algorithms. This also demonstrates the importance of the human factor in platform dynamics, something that is often underestimated by techno-deterministic studies (e.g. Just & Latzer 2016; Klinger & Svensson 2018). Algorithms do not solely dictate user behaviour; they are involved in an interplay with human actors who act on emotions and are not always predictable. In addition, human actors are also present on the platform side in the form of content moderators, programmers and other decision makers that influence how a platform is governed (Gillespie 2018; Roberts 2018).

The most visible human actors on a social network are opinion leaders or "influencers", who are often responsible for bringing information to the attention of their audience. Following the principle of two-step flow communication (Katz & Lazarsfeld 1955), these opinion leaders are powerful because the people they influence are less politically active and informed, therefore being easier to manipulate. When certain political views are repeatedly echoed within one's environment, it is more likely that one will be swayed to believe the same. It is precisely this echoing of opinions in which the social media platforms excel. Studies demonstrate that two-step flow communication functions much the same on social media (e.g. Soffer 2019; Friggeri *et al.* 2014). The aforementioned circulation of content from traditional media to social media means that almost all communication is mediated through individuals (Shirky 2011). These individuals rarely "just" share content; they choose what they want to share (and what they don't) and how they present it. Content circulation on social media is therefore always subject to selection and framing (Wieringa, De Winkel & Lewis 2017). While this is true to a certain extent for all media (see e.g. Hall 1973), social media users do not carry the same responsibilities for independence and objectivity that traditional media do. Information, facts and numbers can therefore be freely manipulated, omitted or framed without any direct punitive consequences. Furthermore, so-called "alternative media" are becoming increasingly more prevalent. These small-scale media productions, which often take the form of blogs, podcasts or videos, present a critical narrative that strays far from the mainstream consensus. They seek to "give voices to the voiceless" and function as a counterinstitution (Fuchs 2010). Their critical approach to the dominant narrative means that the line between alternative media and fake news is diffuse, and these types of media have been seen as a factor in the increasing polarization of the political landscape and normalizing extremism (Spohr 2017; Fletcher 2019).

The topic community surrounding #nexit can be seen as a result of polarization, as it overlaps largely with the far-right audiences on Twitter. We can identify #nexit as one of several topics representing nationalist sentiment. A recent report by the conservative Konrad Adenauer Stiftung (Fischer-Bollin 2020) emphasizes that European nationalism is quite diverse and the references vary in the different member countries. Nationalists across Europe use different topics to trigger the nationalist sentiment of their audiences. In the Netherlands, we can clearly identify far extreme topic communities engaging with topics such as COVID-19 regulations (Veerbeek 2020; Van de Ven and Van Gool 2020), climate change (Hess & Renner 2019; De Kraker *et al.* 2014), the tradition of blackface in Zwarte Piet (Van Es, Van Geenen & Boeschoten 2014; Hilhorst & Hermes 2016), Black Lives Matter, refugees and migration (Mertens, d'Haenens & De Cock 2019;), the government policies to reduce CO_2 emissions (Bakker & Schäfer 2020, n.p.), and the EU. These topics often overlap and connect to each other. The same goes for #nexit. Like many others, this topic is driven by incidents, actualities and its mainstream media coverage. Social media are essential to connecting the various audiences and disseminating messages, and regularly mainstream media amplify these fringe conversations through a disproportionate representation on their own platforms. Studies into neo-nationalism acknowledge the importance of social media platforms for providing the political movement a medium, but too often are they focused on the people and their message, while failing to understand the effect of media practices and platform dynamics (e.g. Postel-Vinay 2017; Udupa 2019). By explicitly studying the effects of platform dynamics on the neo-nationalistic #nexit community, I seek to address this gap in academic knowledge.

In conclusion, an accurate analysis of political communication on social media – or rather, any kind of social media study – should be aware of the many dynamics at play on the platform and beyond. Platforms do not stand on their own; they are part of a large ecosystem consisting of interconnected online services as well as traditional offline media. Speaking of platform dynamics solely in terms of algorithms does not consider the socio-technical complexity of social media, where users, media, businesses, programmers, automated systems and many more actors engage in an opaque web of content and interactions.

Methodology

The central subject of this analysis is a dataset consisting of 28.667 tweets that used the hashtag #nexit in the period March 1st to June 1st 2020. The dataset contains not only the tweets themselves, but also a vast amount of associated Twitter metadata (97 datapoints per tweet in total). To uncover the platform dynamics within this dataset, we must detect patterns within the data. To do so, I build upon digital methods (Rogers 2013) and cultural analytics (Manovich 2009), both of which have been proposed as methods for the socio-technical analysis of rich datasets. Manovich calls for "large-scale computational analysis and interactive visualization of cultural patterns", leveraging the rapidly increasing processing power that is available to scholars to study large datasets in real time. He describes the movement of data as "flows" (Manovich 2016), emphasizing that data is not a static object, but rather exists in constant motion and interaction. A similar approach is put forward by Rogers and the Digital Methods Initiative, who argue for the "fruitfully combining" of digital objects to extract meaning from exploitable and readily available data. Digital methods observes how online devices use data and comes up with ways to repurpose them for social-cultural research (Rogers 2013). Both methods are intentionally broad and non-specific, but rather serve as a framework to build upon. Acknowledging the complexity of modern data and the rapid pace at which it evolves, they understand that explorative and experimental methods are better suited to keep up with the ever changing data landscape. Following digital methods' logic of repurposed data, we can approach this Twitter dataset as a reflection of the platform's inner workings. The datapoints grant us insight into what data the designers use to structure and govern the platform, and what data fuels the algorithms (although we cannot be sure of the full extent of data collection, as noted earlier). Repurposing this data allows us to hermeneutically illume the underlying mechanics.

Some best practices have emerged to help analyze these types of datasets, including programs that have become essential tools for socio-cultural data analysis. I employ two of these for the purpose of this research: the open-source network visualization software Gephi and the statistical programming language R. In the following, I will go into how the dataset has been collected, how Gephi and R are used to analyze the data and the indicators that were determined to detect and define platform dynamics and fake news.

Dataset and data collection

The Twitter dataset that is central to this research was originally collected as part of a social media network analysis that was commissioned by the Dutch representation of the European Commission and conducted by Mirko Tobias Schäfer and me in the context of Utrecht Data School (Bakker & Schäfer 2020, n.p.). The dataset consists of "original" tweets, meaning all tweets excluding *retweets*. Retweets can be seen as direct "copies" of other tweets that are shared on one's own timeline and are useful for analyzing the spread of content on the platform. The lack of these retweets in the dataset means that the focus will not be put on the spread of the tweets, but rather the interconnectedness of their content. The dataset is set up to capture the active debate and its participants, or "explicit participation" (Schäfer, 2019).

The data was requested from Twitter's application programming interface (API). Twitter offers several tiers of APIs, all with different levels of access to the platform's data. Premium APIs and enterprise APIs are offered on a subscription basis, while a free "standard" API is available to developers and researchers.[1] To gain access to the API, one needs to undergo an online screening process, which includes filling out an intent form and a questionnaire. These help the platform determine if the intended use of the data is in line with their privacy conditions, which include a list of restricted use cases: applications of the data that are against the local Internet privacy law, mainly meaning the storing of sensitive information such as health, sexual orientation and political affiliations, as well as the matching of Twitter account information to offline personal details and other social media accounts.[2]

As a researcher at Utrecht Data School, I had access to a standard developer account. To request the data from the Twitter API, I used the programming language R and the open-source package *rtweet*, which is authored by Michael W. Kearney.[3] I wrote a script using functions from *rtweet* that connected to the API and requested tweet data based on a few criteria: the tweets had to be in Dutch, they had to be an "original" tweet (not a retweet) and they had to use the hashtag #nexit. The resulting dataset represents all *active* participants within the #nexit discourse.

Even with the freely available standard API, Twitter offers a wide range of metadata. These datapoints are used by the platform to organize the content on the platform – to recommend potentially interesting content to users or to recognize trending topics, for example – but they are also a commodity that are being sold to commercial partners (Van Dijck, Poell & De Waal 2018; Wieneke & Lehrer 2016; Estrada-Jiménez *et al.* 2019; Ruckenstein & Granroth 2019). Our dataset contains 97 categories of metadata. It is important to note that these datapoints do not represent the full extent of data Twitter collects; some of it is limited to the paid APIs, while more still is a well-kept trading secret of the platform.

Twitter's API limits access on a temporal level as well, allowing users to request tweets going back only nine days. This limitation means that it is difficult to place the dataset in a historical context. I address this restriction by collecting a secondary dataset through the commercial data firm OBI4WAN, which owns an extensive backlog of Twitter data. This dataset is very limited when compared to the dataset that I have collected through the Twitter API, containing only a fraction of the metadata, so it does not lend itself to a comparative analysis. It does however grant me a wider temporal perspective on the #nexit topic community and its popularity. In addition to this, it also helps to bridge two minor temporal gaps in the dataset, which were caused when the R script failed to collect data from the API.

---

[1] https://developer.twitter.com/en/products/twitter-api
[2] https://developer.twitter.com/en/developer-terms/more-on-restricted-use-cases
[3] For *rtweet*'s GitHub repository, see: https://github.com/ropensci/rtweet

Our dataset is limited because of its original purpose, as retweets were deliberately left out of the research design for the European Commission. The decision to not include any retweets was made based on a number of considerations. For the research we conducted for the European Commission, it was decided early on that we were interested in the accounts that shaped the online debate – those that actively participated. The goal of that particular research was to gain insight into how the online debate surrounding the European Commission unfolded and which accounts drove the discourse. To this end, we chose to exclude retweets from the dataset as these do not represent the *active* participants within the online discourse. These tweets represent explicit participation, as opposed to implicit participation (Schäfer 2019), meaning that they add new content instead of spreading it.

Retweet and follow behavior has been the subject of earlier studies of political communities on Twitter, where they were used to measure homophily and the circulation of content within a topic community (see e.g. Conover et al. 2011; Feller et al. 2011; Boutyline & Willer 2015). This research seeks to provide a new perspective by *distant reading* the content that is being produced within the #nexit topic community, and how it relates to Twitter's platform dynamics. Platform dynamics do not only manifest themselves in the way content spreads, but also within the content itself. Through options such as @-mentions, hyperlinks, replies and hashtags, tweets and users are connected within a network of datapoints. This dataset allows us to analyze the dynamics that arise from these connections.

To find and analyze platform dynamics within the dataset, we need to detect patterns. Platform dynamics can take many forms, but they can be broadly divided up into two categories: intraplatform dynamics and media ecosystem dynamics. The former deals with platform dynamics on Twitter itself: how the topic community manifests itself on the platform, whether there is a low or high degree of connectivity between accounts, how it relates to other topic communities, et cetera. The latter, media ecosystem dynamics, deals with the platform's connections to other media in the overall media ecosystem. It should be noted that these categories are not as strictly cut off from each other as this categorization might suggest. Platform dynamics are the result of an interplay of media practices and platform design, meaning that there is a high level of interwovenness between the two categories. That said, the categorization of intraplatform and media ecosystem allows us to better hone in on specific phenomena.

Table 1 shows the indicators in the dataset that can point to platform dynamics. Most of these indicators can be found in the metadata within the dataset itself, often through combining datapoints. Others, such as the percentage of fake news and temporal popularity, require outside information to find correlations with the patterns in the dataset. These indicators were chosen based on the metadata that is available in the dataset and their relationship to the platform dynamics described in the theoretical framework.

Intraplatform dynamics indicators	Media ecosystem dynamics indicators
• Related topic communities • In degree and out degree (user activity and interactivity) • Clusters of connected accounts • Hashtags used in conjunction with #nexit (indicators of related themes) • Influence: number of followers, activity (number of tweets) • Explorative analysis of images	• Most shared hyperlinks / domains • @-mentions and network position of accounts of traditional news outlets • Percentage of alternative media and fake news sites • Temporal popularity: activity surrounding news events • Explorative analysis of images (depictions of media)

Table 1. *Data indicators for platform dynamics in intraplatform and media ecosystem context.*

To analyze these indicators, I make use of Gephi and R. I will briefly reflect on how both are employed.

Gephi

Gephi allows us to visually explore the network, in addition to detecting clusters and analyzing patterns through algorithms. The network is represented through nodes and edges. Connections between nodes are determined on based on three types of Twitter interactions: @-mentions, replies and quotes. The number of appearances of an account within the dataset determines the size of that node in the network. This means that the size of a node increases when there are a lot of interactions with that particular account, but also when the account itself is actively tweeting. All network visualizations in this paper have been processed with the sorting algorithm ForceAtlas2 (Ventuini, Heymann & Bastian 2014), which spatializes nodes based on their prominence in the network. The Leiden algorithm (Traag, Waltman & Van Eck 2019) has been used to detect communities. These, in turn, have been colorized with randomly generated colors.

Two types of networks are studied in this paper. The first only contains accounts and the connections between them, as described above. The second contains these accounts, as well as the hashtags. Hashtags play an important role in Twitter networks, linking content and serving as anchor points for topic communities (Wieringa *et al.* 2016). Incorporating hashtags in the analysis allows us to find related topic communities and themes. Using the findings of these analyses, we can drill down on specific phenomena for further research.

It should be noted that Gephi allows for many different types of visualizations, all of which can potentially lead to a different perspective on the network. The visualization methods used for this analysis have been discussed extensively with other experienced researchers at Utrecht Data School, and have previously been successfully applied in my research for the European Commission.

R

In addition to retrieving data from the API, R is used for cleaning the data and performing statistical analysis and *distant reading* on the dataset. This includes extracting domain names from the hyperlinks to find the most commonly referenced sources in the dataset, plotting bar charts and formatting the data for analysis through Gephi. I used RStudio, a more user-friendly interface for R. As mentioned above, I utilized the package *rtweet* for data retrieval. Other packages I used include *httpuv* for the handling of HTTP requests[4]; *knitr* for dynamic report generation[5]; *lubridate* to handle date and time data[6]; and *tidyr* to clean up the large dataset.[7]

Defining alternative media

The relation between the #nexit topic community and the overarching media ecosystem will be explored through a domain analysis, which looks for the most popular hyperlinked sources in the community. These domains will be held against the fake news index *Hoax-wijzer*, to find how the topic community relates to alternative media and hoax sites.[8] For websites that are not included on the *Hoax-wijzer*, I follow the fake news detection methods defined in a literature survey by Zhou and Zafarani (2020), who distinguish eight types of fake news based on authenticity, intention and whether the information is news. I shall refer to these concepts collectively as "alternative media". The sites that fall under this umbrella term all push a certain agenda by misconstruing information or constructing false stories. Through this, they push back against the narrative put forward by mainstream media.

Concept	Authenticity	Intention	News?
Deceptive news	Non-factual	Mislead	Yes
False news	Non-factual	Undefined	Yes
Satire news	Non-unified[2]	Entertain	Yes
Disinformation	Non-factual	Mislead	Undefined
Misinformation	Non-factual	Undefined	Undefined
Cherry-picking	Commonly factual	Mislead	Undefined
Clickbait	Undefined	Mislead	Undefined
Rumor	Undefined	Undefined	Undefined

Table 2. *Comparison between concepts related to fake news (Zhou and Zafarani 2020).*

[4] https://cran.r-project.org/web/packages/httpuv/index.html
[5] https://cran.r-project.org/web/packages/knitr/index.html
[6] https://cran.r-project.org/web/packages/lubridate/vignettes/lubridate.html
[7] https://cran.r-project.org/web/packages/tidyr/index.html
[8] http://www.hoax-wijzer.be/

United through opposition: analysis and findings

In this chapter, I present the findings from my analysis of the #nexit dataset. We start off by looking at how the community establishes itself on Twitter through network and hashtag analysis, as well as an explorative study of the images shared within the cluster. We then zoom out to look at the position of #nexit in the larger media ecosystem and how the cluster relates to news events and actualities.

Tightly connected, loosely defined

The dataset consists of 28.667 Dutch tweets that have used the hashtag #nexit in the period March 1st to June 1st 2020. When visualized and analyzed in Gephi, the network that appears shows a remarkable level of interconnectedness. Figure 1 shows the *active* accounts in the dataset, who have posted at least two tweets containing #nexit in the analysis period: a total of 1.631 accounts. The average degree is 4.7, meaning that on average a single account is connected to 4.7 other accounts within the network. This high level of interconnectedness also shows in the cluster analysis: 40% of the actively tweeting accounts in the network are part of the same cluster. Even more remarkable is the fact that there are no alternative clusters to speak of, as the remaining 60% are mostly disconnected from each other and the main cluster. This indicates that no other groups on Twitter use the hashtag. The high level of interconnectedness and the lack of opposing voices point to an ideological echo chamber that reinforces a particular worldview (Wieringa *et al.* 2016). Users within this echo chamber interact with many of the same hashtags, accounts and content, meaning that their algorithmic profiles will share many similarities, which could in turn lead to further homogenization through content recommendations. We shall see further proof of this as we continue.

The term "Nexit" refers to the scenario of the Netherlands leaving the European Union, similar to the British Brexit, which had become a popular term leading up to the United Kingdom's departure in 2020. Nexit supporters give a myriad of reasons for this departure: there would be not enough financial benefits in it for the Netherlands; EU budget is used to support unstable and poorer economies; the European Union pushes an agenda of monoculture that is unrealistic, et cetera. These claims do not have widespread support in the Netherlands; according to the November 2019 Eurobarometer, 22% of Dutch citizens feel negative towards the EU, of which only 3% responded with "very negative". In contrast, 43% responded positively.[9]

Through our analysis of this dataset, we can conclude that the users who actively engage with the #nexit topic community represent the small minority of Dutch citizens who display a very negative stance towards the European Union. Within this community, the term "Nexit" does not so much refer to a political issue, but rather an expression of a broader anti-establishment sentiment. The hashtags that are used in conjunction with #nexit speak volumes: #stemzeweg ("vote them down" – 371 tweets), #EUSSR (comparing the European

[9] *Eurobarometer November 2019. In general, does the European Union conjure up for you a very positive, fairly positive, neutral, fairly negative or very negative image?*
https://ec.europa.eu/commfrontoffice/publicopinion/index.cfm/Chart/getChart/themeKy/19/groupKy/102

Union to Sovjet Union – 274 tweets) and #kominverzet ("start the resistance" – 199 tweets). These hashtags do not refer to specific issues related to the EU, but rather a general anti-authoritarian attitude. The extremity of the statements could be the result of a trend of radicalized Euroscepticism, as noted by Galpin (2017). She notes that since the UK's departure from the EU, Euroscepticism in mainstream media and the social media sphere has become much more prominent. Negativity has historically proven to gain more attention from the public – which is not necessarily unwelcome, as it subjects those in positions of power to a healthy level of scrutiny. When negativity is rewarded in the form of reader attention, however, it is a matter of time before it slips into a "spiral of cynicism" (Galpin 2017). With social media leading what has become widely recognized as the "attention economy" (Davenport & Beck 2011), where keeping users engaged as long as possible is key, negativity has proven to be a valuable tool. Negative emotions such as anger, fear or frustration seem to drive many posts on Twitter. A number of publications have pointed out that the recommendation algorithms seem to favor this kind of content (Park 2015; Vargo & Hopp 2020; McCreery & Krach 2018). Emotionalizing content draws more engagement and is seen as fitting the inherent purpose of the social media platform, namely to engage and connect users.

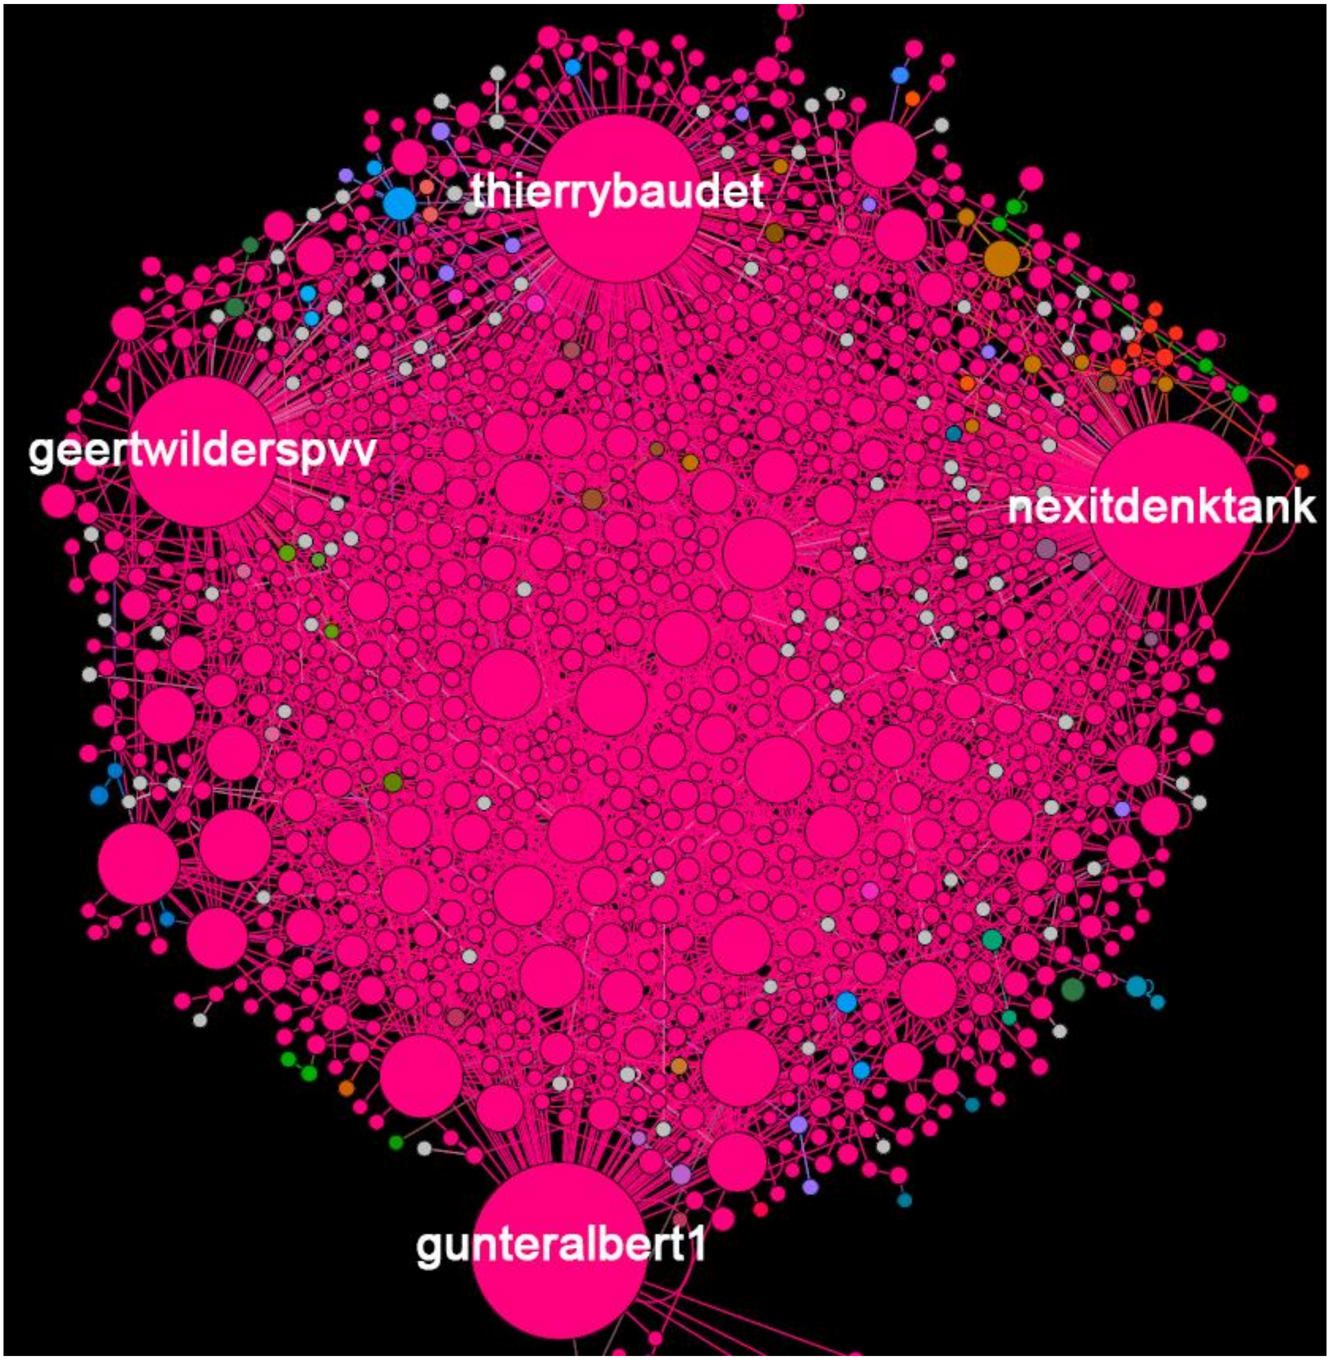

Figure 1. *Network visualization of the active participants in the #nexit cluster. Shown are all accounts who have posted two or more tweets containing the hashtag #nexit. The size of the nodes is determined by the number of occurrences of the account within the dataset.*

Figure 1 shows the network visualization of all accounts with an out-degree of 2, i.e. accounts that have posted two or more #nexit tweets. We can distinguish four prolific active accounts: @geertwilderspvv (Geert Wilders – 820K followers), @thierrybaudet (Thierry Baudet – 230K followers), @nexitdenktank (official account for anti-EU action group and alternative news site nexitdenktank.nl – 10,9K followers) and @gunteralbert1 (anonymous far-right account - 133 followers). These accounts all represent a different character of the cluster, and they all play a different part in constructing an ideological echo chamber. The first two are run by popular populist politicians, @nexitdenktank is the account of an anti-EU organization that also posts blogs on its website, and the latter is an extremely active but ultimately non-influential anonymous account, who has posted 668(!) #nexit tweets in the analysis period. Politicians relay the #nexit message to the traditional media, and through this lend legitimacy and exposure to the term. Accounts such as @nexitdenktank function as opinion leaders: their large follower count and self-curated presentation as a spokesperson for #nexit community lead to visibility and legitimacy, and their website provides the #nexit users with arguments and sources. An account such as @gunteralbert1 is an extreme version of a type of user that is very visible within the #nexit cluster: the "one issue account". These accounts are focused on a single subject and relate practically all information back to this topic. While not popular, these accounts add to the prominence of the community on the platform through sheer volume: their high level of interactivity and use of the #nexit hashtag play into Twitter's trend discovery algorithms, and thus increase the chance of #nexit appearing in the trending topics (Scott 2015: 10).

Apart from the aforementioned accounts, there is a number of accounts who do not actively participate in the #nexit discourse, but are still visibly part of the network (Figure 2). These accounts are integrated in the network because they are @-mentioned or cited by active accounts. @MinPres (the official account of Dutch prime minister Mark Rutte) is by far the most prominent node within the network; other important political figures such as @RobJetten (parliamentary leader for the social-liberal D66) and @TimmermansEU (Frans Timmermans, European Commissioner and vice-president of the European Commission, member of the social-democratic Partij voor de Arbeid) are also frequently mentioned in the context of #nexit. These three politicians are referenced through their personal accounts, but D66 is the only one of their political parties whose Twitter handle is frequently mentioned. These figures and this party function as a scapegoat for all EU-related issues within the #nexit community, as will become clear in the following chapter where we look into the visual identity of the cluster.

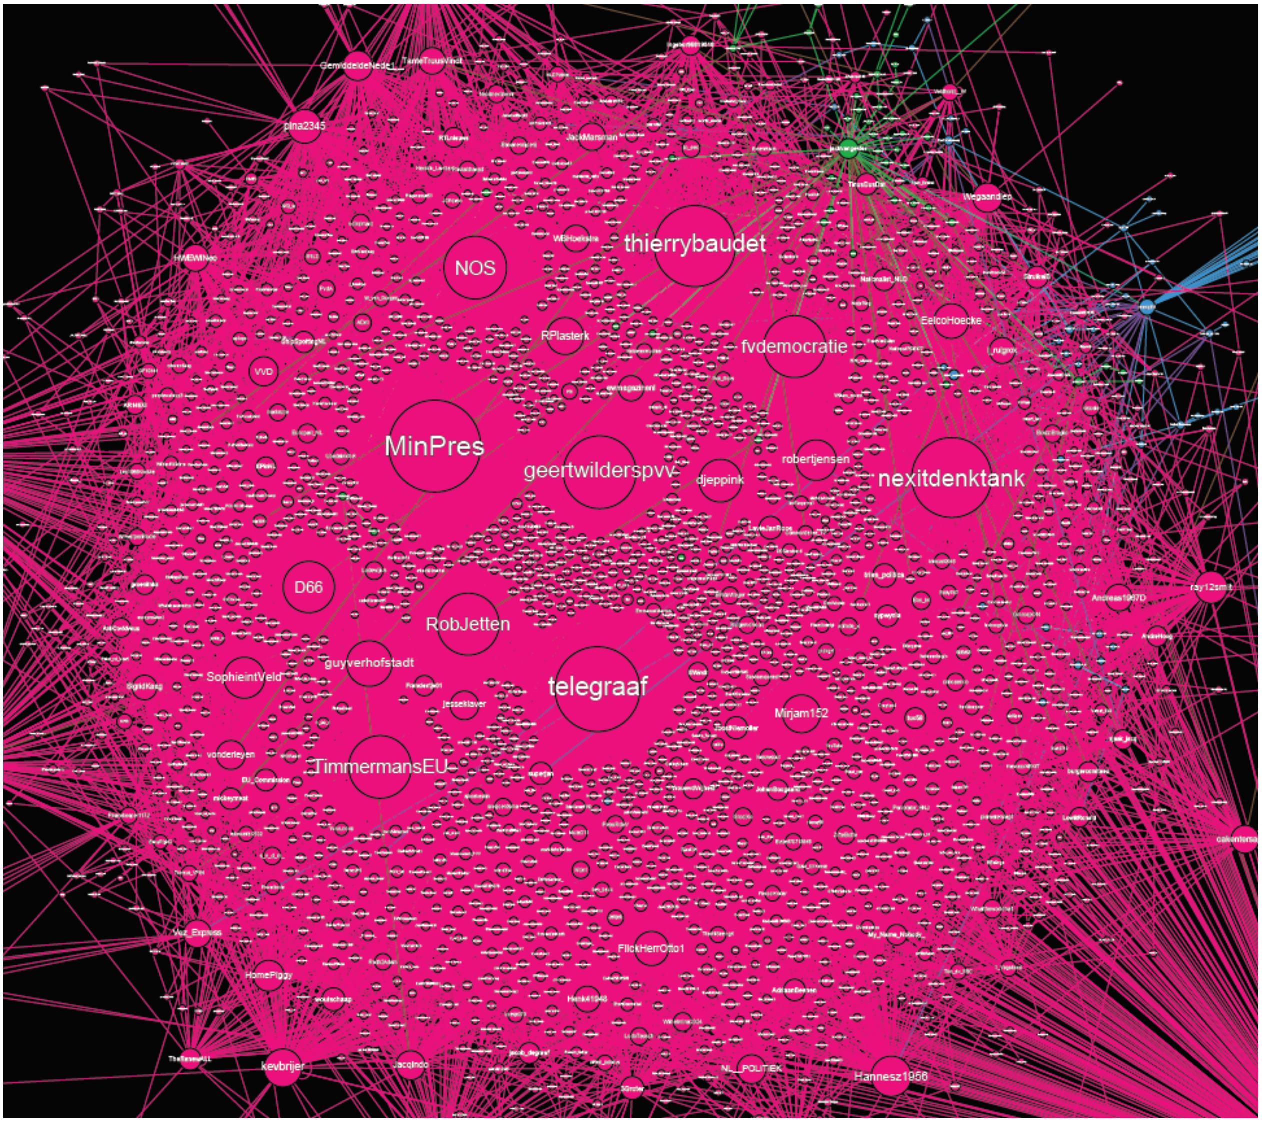

Figure 2. *Network visualization of the #nexit cluster, including non-participating accounts. The size of the nodes is determined by the number of occurrences of the account within the dataset.*

Wrecking balls and swastikas: #nexit's visual identity

A picture is worth a thousand words, and this adage is certainly true in the #nexit cluster. While only 20.6% of the tweets within the dataset contain some sort of visual element in the form of an attached image or a hyperlink preview, 76.6% of the of the 500 most retweeted tweets include visual media, indicating that the use of visuals is highly effective on Twitter.[10] An explorative analysis of these images reveals a distinct visual identity, as demonstrated in Figure 3. We can note two clear subjects in these images: flags and politicians.

The dominant icon in these images is the European flag, with its blue backdrop and yellow stars showing up in the vast majority of images. Rarely is the flag used "as is". Elements such as swastikas, a hammer and sickle, fire or handcuffs are added to alter its political meaning and call to mind oppressive regimes. These images show similarities to political cartoons in wartime; a similarity that is further enhanced by the aforementioned "war-like" hashtags (e.g. "start the resistance"). Inversely, the Dutch flag is used as an ideal or a symbol of power, for example as a wrecking ball destroying the European flag.

The number of politicians who are represented in these images is limited, with Dutch prime minister Mark Rutte and European commissioner Frans Timmermans making up the majority of faces in the posts. Similar to the images of the European flag, their visage is not safe from photo editing. Rutte is pictured throwing money in a toilet, growing an elongated nose from constant lying, or wearing diapers, amongst others. Made up quotes along the lines of "I'm giving all your money to the EU and there is nothing you can do about it" are added into pictures of Rutte laughing hysterically. These images present a worldview in which a small few politicians can single-handedly control a country's future. Rutte and Timmermans function as scapegoats here: they are responsible for all issues the #nexit cluster concerns itself with. This worldview fails the acknowledge the complexity of politics, and seems to take cues from the caricatural way United States president Donald Trump presents his political opponents – something that is further corroborated by Trump's presence in many images and Twitter bios in the #nexit cluster.

On a platform where conciseness is a necessity, a striking image is an effective method of conveying a message (Coesemans & De Cock 2017). The #nexit community has seemingly agreed upon a visual language that is both immediately recognizable and capable of propagating a particular strand of Euroscepticism. Through extreme exaggeration, drawing upon historical stereotypes of corruption and evil, the posters of these images create a dystopian representation of the European Union that plays into the fear and anger that drive engagement on social media. The effectiveness of this method is demonstrated by the fact that the vast majority of popular tweets within the dataset contain a visual element.

---

[10] Twitter recommends the use of visuals to businesses on their blog, demonstrating through their data that photos and videos fuel more engagement: https://blog.twitter.com/en_us/a/2014/what-fuels-a-tweets-engagement.html

Figure 3. *A selection of images demonstrating the visual language of #nexit. The full dataset contains 1.315 images.*

#nexit within the media ecosystem

Twitter does not stand on its own – it exists in a web of connections within the media ecosystem. For many users, social media networks bring the content they are interested in together in one place and function as a portal to the rest of the Internet (Fletcher 2019). In this chapter, we increase our scope and look beyond Twitter to see how the #nexit cluster positions itself within this larger media ecosystem. We start off by analyzing the hyperlinks that are shared in the #nexit tweets. Then, we focus on the sharing of media through images. Finally, we analyze how activity within the community is tied to external news events.

Communication on Twitter is not limited to the text and media in the tweets themselves. Often, the platform serves as a gateway to external sources, referring users to websites beyond Twitter's control. In the dataset, 5.109 tweets, or 17%, contain a hyperlink, though a massive 3.558 of those link "internally" to other pages or tweets on Twitter. The remaining 1.551 hyperlinks, or just 5.4%, refer to external websites. Despite these small percentages, tweets that contain hyperlinks draw far more engagement, with 272 out of the top 500 most favorited and retweeted tweets containing a hyperlink. The reach of these links is therefore larger than the percentage of hyperlinks might suggest. This seems to tie into the higher level of engagement with visual content that we noted earlier, as hyperlinks on Twitter are formatted to automatically include an image from the target website, thus making the tweet stand out among "plain" text tweets.

To understand how the #nexit community takes shape on Twitter, we must look beyond the platform and see which sources are commonly linked in these tweets. Figure 4 shows the 30 most hyperlinked domains in the dataset, excluding the internal Twitter links. We can note a notable presence of mainstream media within the cluster: the mainstream news outlets De Telegraaf (268 links), NOS (160 links), AD (93 links) and NU.nl (46 links) take up 4 spots in the top 5 most frequently hyperlinked domains (Figure 4). De Telegraaf, a traditionally populist newspaper that has started to lean more far-right in recent years (Elibol, Hagen & van de Ven 2019), visibly dominates the news domain, which could have to do with it being the most popular mainstream outlet that aligns with the political identity of #nexit. Interestingly, only the accounts of NOS and De Telegraaf are prominently featured in the network (see Figure 2), both for different reasons. The public – and therefore state-funded – NOS is not the most popular Dutch news outlet on Twitter (that honor goes to NU.nl with 1.5 million followers, as opposed to NOS' 1.1 million), but it does represent the establishment that the #nexit community explicitly opposes. On the other hand, De Telegraaf aligns itself more with the Eurosceptics and does not scare away from using sensationalistic headlines – something that plays into the lack of nuance within the #nexit community. Both of these outlets are often referenced by accounts within the #nexit network because they represent a different side of the debate.

The video streaming platform YouTube is the second most hyperlinked website in the #nexit cluster. YouTube, which allows anyone to upload video content, has become a popular vehicle for spreading fake news and conspiracy theories in recent years (Hussain *et al.* 2018; Heydari *et al.* 2019). An explorative analysis of the links in our dataset shows that the vast majority of the YouTube hyperlinks lead to videos containing misinformation. Alternative media productions such as CommonSenseTV, Ongehoord Nederland and Jensen, all known outlets for conspiracy theories, are especially well represented. These channels also have a strong Twitter presence, using the platform to cross-promote their different social media accounts. While their Twitter accounts do not show a high level of engagement in our network analysis (Figure 2), they do rank among the active accounts with the most followers in the dataset, suggesting that their popularity on Twitter is a result of activity across different platforms.

YouTube content on Twitter has another major effect on platform dynamics. When a YouTube video is shared on Twitter, it is embedded within the tweet, meaning that the user can watch the video without leaving Twitter. YouTube is a subsidiary of Google LLC and thus collects user data for the technology company that, much like Twitter, uses this information to build a data profile and recommend more similar content to its users. A number of studies have criticized YouTube for leading users into "rabbit holes": a pipeline of recommended content that becomes gradually more radical, often leading to fringe far-right ideologies (e.g. Kaiser & Rauchfleisch 2019; Ribeiro *et al.* 2020; Hussein, Juneja & Mitra 2020). If a user becomes active within the #nexit community, there is a considerable chance that they will interact with the abovementioned YouTube content. Their political interests will then be recorded in their data profile of both Twitter and Google, extending the reach of the ideological echo chamber across multiple platforms. Similar dynamics are at play when news outlets or blogs embed social media content in their websites, which allows tracker cookies from social networks to gain access to the user's behavior across different sites (Lerner *et al.* 2016; Mehrnezhad 2020).

Moving further into the realm of alternative and fake news, we find that 12 out of the 30 most frequently shared URLs belong to this category (Figure 4). Six of these are blogs that are included on the *Hoax-wijzer*, while the others adhere to the previously discussed alternative and fake news categories (Zhou and Zafarani 2020). These sites all lean far-right, and share a myriad of conspiracy theories and misinformation. Many are aimed specifically at peddling a Eurosceptic message. Interestingly, links to these sites prove to be far less common than references to mainstream media. Fletcher (2019) establishes that online news outlets in Western Europe generally have mixed audiences, with much overlap in news consumption over the political spectrum. Those leaning far-right or far-left tend to engage more with fringe publications, but also come into contact with mainstream media. These findings are confirmed by the combination of dominant mainstream media and far-right blogs in the #nexit cluster.

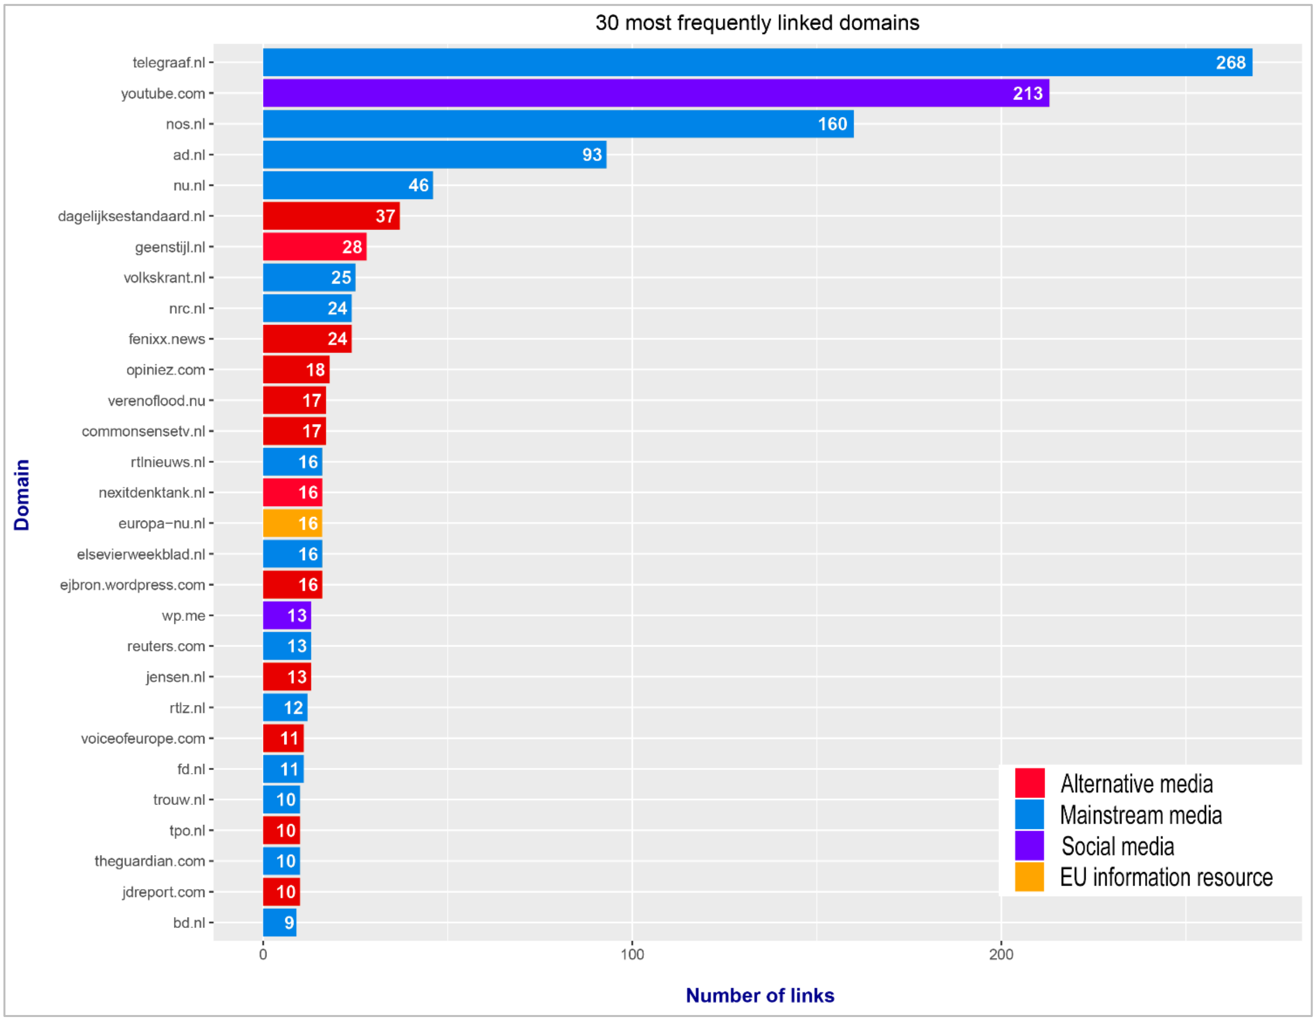

Figure 4. *30 most frequently shared URLs in conjunction with #nexit, categorized by type of website. The category "alternative media" is defined as websites that are included in the* Hoax-wijzer, *or that adhere to one or more definitions put forward by Zhou and Zafarani (2020).*

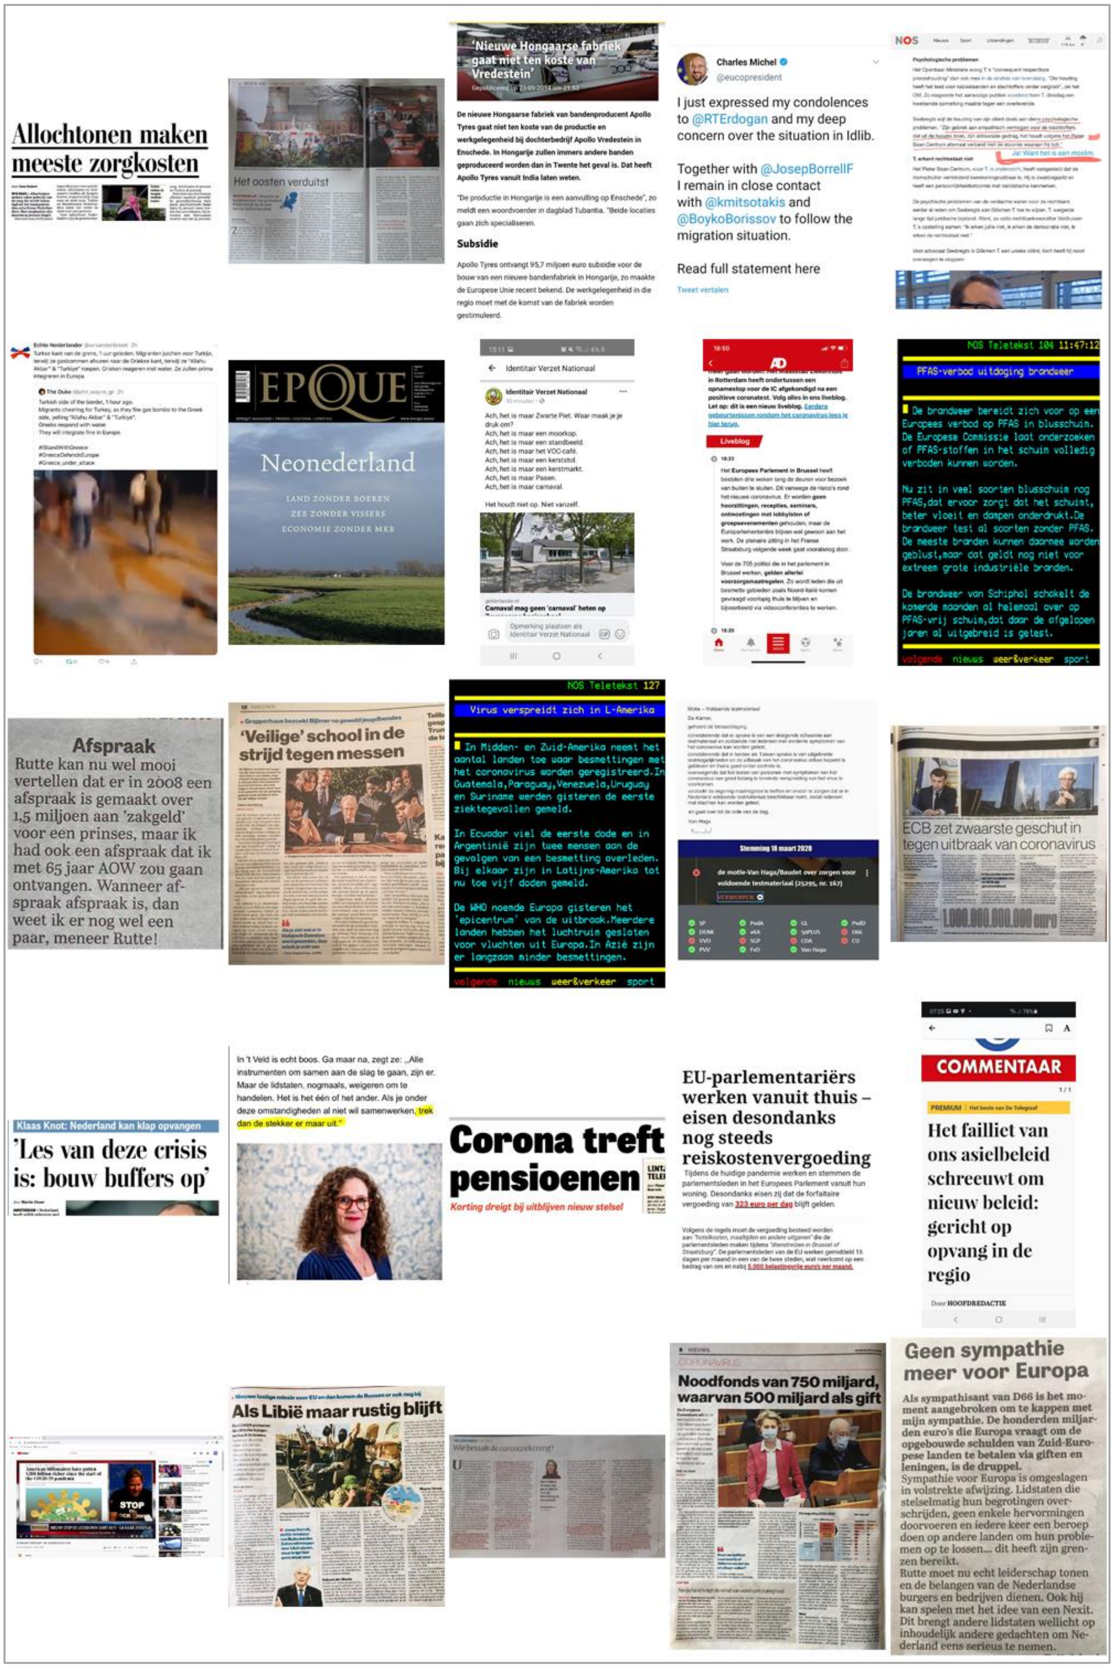

Figure 5. *A selection of media clippings shared within the #nexit cluster.*

Hyperlinks are not the only way media can be shared. Figure 5 demonstrates the different ways in which media are shared through images. Offline media, such as newspapers, must be mediated to be shared in the online sphere, through photographs, scans or transcriptions. This mediation has major consequences for framing. Figure 5 shows a clear preference for headlines. Often the article is cropped so that only the title or a small part of the text is visible, and if the full article is photographed, the image resolution makes it near impossible to make out text other than the headline. Headlines are inherently suggestive; they are composed specifically to tempt a reader, often by focusing on the most gripping part of a story. By selectively focusing on these headlines or small fragments of text, the meaning of the article is altered and nuance is lost (Blom & Hansen 2015; Kuiken *et al.* 2017). A similar type of framing can be found in the underlining or circling of specific parts of the article, which puts emphasis on small fragments of text and distracts from the rest of the article.

These framing devices demonstrate that the principle of two-step flow communication (Katz & Lazarsfeld 1955) is inherent to the mediation of offline media to the Internet. The person sharing the content not only dictates how that content is framed, they also take away the option for the recipient to directly access the text themselves – something that *is* possible through hyperlinks. Further framing occurs in the tweet itself, which accompanies the image with commentary and the #nexit hashtag. In this way, media outside of Twitter is taken out of context and weaponized to support the neo-nationalistic #nexit message.

Just like fake news, altered media is also rampant on social media. Screenshots can be edited or faked to suggest false information or accuse people of saying offensive things (Shen *et al.* 2018). Close reading of these sources could reveal such practices, though this unfortunately is not manageable within the scope of this research.

Peaks and valleys: temporal analysis

The popularity of topic communities is strongly tied to news and actualities. Plotting the use of #nexit (Figure 6 & 7) shows a distinct pattern, with periods of relative calm alternated by sudden spikes in popularity, often more than doubling the use of the hashtag on a given day. Before we analyze the temporal data, some context is needed. The analysis period was an eventful time, as the first COVID-19 cases started appearing in the Netherlands just two weeks after data collection had commenced. The corona pandemic led to lockdowns and citizens looking towards their national governments for regulations and information. Suddenly, the borders between European countries became tangible. In terms of corona regulations the European Union stood relatively powerless, as individual member states decided upon their own course of action. The European Commission did however propose a number of relief funds for the corona-stricken Southern and Eastern Europe. Considering this, let us turn our attention to the pattern in Figure 6 and 7.

The historical graph in Figure 6 visualizes the use of the hashtag #nexit on a day-to-day in the period March 1st to June 1st 2020. This graph *includes* retweets. Three prominent peaks jump out: March 31st (3.862

tweets), April 21st (2.571 tweets) and May 27th (3.259 tweets). Between these spikes in usage, the hashtag normally lingers in the 700-1500 tweets per day area. Figure 7, which *excludes* retweets, shows a similar pattern, with some exceptions. In the following, we will briefly look into the events that trigger the sudden use of #nexit:

31st of March: Prominent Italian figures publish a letter in the German newspaper *Frankfurter Allgemeine Zeitung*, in which they blame the Netherlands for putting the European Union at risk by obstructing the financial support negotiations for COVID-stricken countries. Dutch mainstream media run headlines about the Italian accusation, and outrage is sparked in the #nexit cluster. 939 #nexit tweets are posted (Figure 7), a number that is quadrupled if retweets are included (Figure 6). The number of #nexit tweets dies down the next day, returning to a third of the day before.

21st of April: The European Commission reveals the allocation of money for the Coronavirus Response Investment Initiative, which shows that the Netherlands will only receive 25 million euro from the EU relief fund, as opposed to Poland's 7.435 million and Hungary's 5.603 million, amongst others. @geertwilderspvv shares the infographic on his Twitter page accompanied with "#NEXIT". The tweet gains traction and is retweeted 746 times and gets 1.8K likes, causing a spike in Figure 6. Figure 7, which excludes retweets, does not show a significant jump in absolute usage, however.

27th of May: The European Commission proposes a 750 billion corona relief fund. @geertwilderspvv posts five tweets that contain #nexit, @thierrybaudet also posts one. The use of #nexit increases tenfold, but again dies down just a day later.

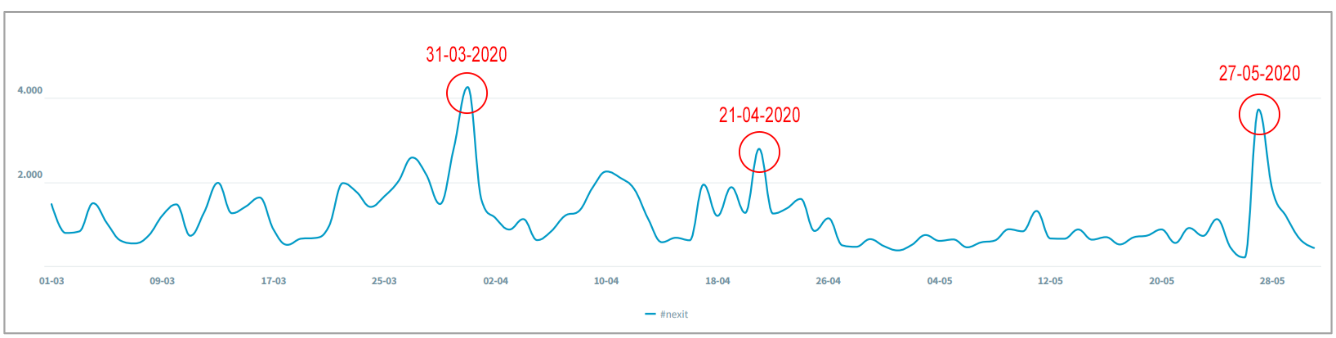

Figure 6. *Historical graph of the popularity of #nexit between March 1st 2020 and June 1st 2020,* ***including retweets****. Data for this visualization is provided by OBI4WAN.*

Figure 7. *Historical graph of the popularity of #nexit between March 1st 2020 and June 1st 2020,* ***excluding retweets****. Data for this visualization is provided by OBI4WAN.*

These temporal graphs supports the observation that neo-nationalism is reactionary in nature (Eger & Valdez 2014; Antonsich 2017; Bergmann 2020). All large spikes of activity in the #nexit cluster are preceded by an external event related to the European Union. In the analysis period, the #nexit community never initiates activity by itself. Twitter as a platform lends itself well to this type of reactionary content, as its character limit necessitates short and "snappy" messages, but is not suited to more nuanced content.

A hashtag such as #nexit can be used to express a general sentiment without the need to further clarify it, as it assumes the reader is aware of what it signifies. On the other hand, this implicitness can lead to a wide variety of interpretations. Throughout this analysis, we have noted time and again that #nexit does not have one explicit meaning, but rather serves to express a broad anti-authoritarian sentiment. The large presence of fake news and seemingly unrelated conspiracy theories in text, imagery and sources, combined with an unconstructive, reactionary tone points to a group that is united through opposition. Twitter allows these users to vent and surround themselves with like-minded people that share their frustration. The large presence of politicians and mainstream media on the platform grants the #nexit cluster the opportunity to agitate "directly" towards those they oppose, creating the illusion of active resistance – even though our temporal analysis demonstrates that this opposition is short-lived, and our network analysis shows that their #nexit message does not reach beyond the ideological echo chamber.

Perhaps the most salient conclusion we can draw from this temporal analysis is the inflammatory role that populist politicians Baudet and Wilders play within the network. Just like the rest of the users, the two only *react to* events with the hashtag, rather than using it as a way of signaling their anti-EU stance in their day-to-day communication. This suggests that they deliberately play into this ideological hub when it suits them, and that they utilize the aforementioned connotations of the term "nexit" to further ignite anti-authoritarian sentiments within the most radical part of their base. This practice of coded communication is known as dog whistle politics, a tactic that is often employed to communicate radical or controversial ideas through a form of disguised discourse that is only understood by a particular group. Donald J. Trump – who, as we noted before, is highly present in the #nexit cluster – has also been shown to employ such dog whistles during his (run for) presidency (Terrill 2017; Cantalupo 2019). My data suggest that Wilders and Baudet make deliberate use of the same covert communication tactics to rally the more extreme part of their base without explicitly expressing radical terms. This impression is further corroborated by the fact that both their political parties, PVV and FvD, are named 225 and 240 times respectively in the profile biographies of the 2.831 #nexit users, showing that 10% explicitly affiliate themselves with one or both of these parties.

Towards a *collective* platform dynamics awareness

Throughout my analysis, the #nexit cluster has revealed itself to be a relatively small, but tightly connected neo-nationalistic community that rallies around EU-related events, but propagates a more diffuse anti-authoritarian attitude that is not always related to the European Union. Rather than functioning as a place for constructive debate, Twitter is used by this cluster as an outlet for frustration. The platform lends itself well to neo-nationalism's reactionary tendencies, which is demonstrated by the brief activity spikes surrounding EU-related news. The socio-technical platform dynamics that arise from my dataset demonstrate that human actors, platform design and the overarching media ecosystem all have a big impact on how online ideological communities are shaped. If there is one thing to take away from this research, then, it is the fact that we must collectively move away from techno-determinism in our approach to the rise of extreme ideologies on social media platforms. Even as recently as this year, the highly popular documentary *The Social Dilemma* (Rhodes & Orlowski 2020) presented its mainstream audience with a representation of social media that echoes Pariser's (2011) flawed notion of the algorithmic filter bubble – demonstrating that the concept is still clearly present in the collective consciousness. This way of thinking hinders effective action against online extremism.

The findings of this explorative research can prove valuable for media scholars, policy makers and designers of social media platforms alike. Media scholars will find steps towards an analytical framework that incorporates the effects of the overarching media ecosystem in the study of platform dynamics and their influence on the formation of extreme ideologies. They could build upon the several types of analysis I put forward in this paper to further investigate these platform dynamics and potentially uncover new ones. Some possible avenues to pursue include comparative studies of dynamics between different platforms, the relation between mainstream media coverage of political hashtags and growth of their userbases, and connections and overlaps between different ideological topic communities. For policy makers, the methods described in this research will help to more effectively determine how and where potentially harmful ideological communities might form, and to build policies that prevent further growth of those groups and dissemination of their ideas. Too often governments take a techno-deterministic approach to policy, putting the blame squarely on social media platforms and their moderation (Schäfer 2019). I have shown that platforms are just one part of a larger whole, and that media and socio-cultural factors are just as crucial. Effective policy must address the socio-technical platform dynamics that go beyond social media, for instance the role of public figures who legitimize extreme sentiments through ideological dog whistles. Ideally, studies like this would also inform platform designers and engineers when considering new platform features and moderation guidelines. The recent measures against conspiracy theories and misinformation on Twitter and other social media platforms[11] are symbolic gestures that do not address the underlying problem: their revenue models reward content that drives user engagement, which too often stems from negativity and anger. Reactionary content therefore performs especially well, as we have seen in the #nexit community. Real, substantial change must also be supported by those who operate the platforms that structure our everyday information flows.

[11] YouTube, Facebook and Twitter all took individual action against the spread of misinformation and conspiracy theories in 2020, mostly focusing on specific conspiracy groups such as QAnon.

Bibliography

Alini, A., A. Datta, J. Li & D. Huynh. (2017, November). *Empirical Analysis of Factors Influencing Twitter Hashtag Recommendation on Detected Communities*. Paper presented at the 13th International Advanced Data Mining and Applications Conference, Singapore. https://doi.org/10.1007/978-3-319-69179-4_9

Alizadeh, M. et al. (2019). Psychology and morality of political extremists: evidence from Twitter language analysis of alt-right and Antifa. *EPJ Data Science*, *8*(17). https://doi.org/10.1140/epjds/s13688-019-0193-9

Antonsich, M. (2017). The Return Of The Nation: When Neo-Nationalism Becomes Mainstream. *Society and Space*. Retrieved from https://www.societyandspace.org/articles/the-return-of-the-nation-when-neo-nationalism-becomes-mainstream

Badawy, A., A. Addawood, K. Lerman & E. Ferrara. (2019). Characterizing the 2016 Russian IRA infuence campaign. *Social Network Analysis and Mining*, *9*(31). https://doi.org/10.1007/s13278-019-0578-6

Bakker, J.J. & M.T. Schäfer. (2020). *Mapped out: Network research into the Twitter communication of the Dutch representation of the European Commission.* Utrecht: Utrecht Data School.

Becker, M. (2015). *Ethiek van de digitale media.* Uitgeverij Boom.

Beigi, G. (2018). *Social Media and User Privacy*. Retrieved from https://arxiv.org/abs/1806.09786

Bergmann, E. (2020). *Neo-Nationalism: The Rise of Nativist Populism*. Springer Nature.

Blackwell, D., C. Leaman, R. Tramposch, C. Osborne & M. Liss. (2017). Extraversion, neuroticism, attachment style and fear of missing out as predictors of social media use and addiction. *Personality and Individual Differences*, *116*, 69-72. https://doi.org/10.1016/j.paid.2017.04.039

Blom, J.N. & K.R. Hansen. (2015). Click bait: Forward-reference as lure in online news headlines. *Journal of Pragmatics, 76*, 87-100. https://doi.org/10.1016/j.pragma.2014.11.010

Boutyline, A., & Willer, R. (2016). The Social Structure of Political Echo Chambers: Variation in Ideological Homophily in Online Networks. *Political Psychology*, *38*(3), 551–569. https://doi.org/10.1111/pops.12337.

Bossetta, M. (2018). The Digital Architectures of Social Media: Comparing Political Campaigning on Facebook, Twitter, Instagram, and Snapchat in the 2016 U.S. Election. *Journalism & Mass Communication Quarterly*, *95*(2), 471-496. https://doi.org/10.1177/1077699018763307

boyd, d. (2010). Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications. In Z. Papacharissi (Ed.), *Networked Self: Identity, Community, and Culture on Social Network Sites* (pp. 39-58). Routledge.

boyd, d. & Crawford, K. (2012). Critical Questions For Big Data. Information. *Communication & Society*, *15*(5), 662-679. https://doi.org/10.1080/1369118X.2012.678878

Bucher, T. (2018). *If...Then: Algorithmic Power and Politics*. Oxford Scholarship Online. https://doi.org/10.1093/oso/9780190493028.003.0003

Bucher, T. & A. Helmond. (2016). The Affordances of Social Media. In J. Burgess, T. Poell & A. Marwick (Eds.), *The SAGE Handbook of Social Media* (pp. 233-253). SAGE Publications Ltd.

Cantalupo, N.C. (2019). Dog Whistles and Beachheads: The Trump Administration, Sexual Violence, and Student Discipline in Education. *Wake Forest Law Review*, *54*(2), 303-362. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3323432

Coesemans, R. & B. De Cock, (2017). Self-reference by politicians on Twitter: Strategies to adapt to 140 characters. *Journal of Pragmatics, 116*, 37-50. https://doi.org/10.1016/j.pragma.2016.12.005

Conover, M. D., Ratkiewicz, J., Francisco, M., Goncalves, B., Menczer, F., & Flammini, A. (2011, July). *Political Polarization on Twitter*. Paper presented at the Fifth International AAAI Conference on Weblogs and Social Media. Retrieved from https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2847

Dauber, C.E. & K. Ilter. (2019). The Relationship between Social Media and Radicalization. In J.R. Vacca (Ed.), *Online Terrorist Propaganda, Recruitment, and Radicalization* (pp. 47-63). Taylor & Francis LTD.

Davenport, Thomas H., John C. Beck. (2001). *The Attention Economy: Understanding the New Currency of Business*. Harvard Business Press.

Dijk, Van J., T. Poell & M. De Waal. (2018). *The Platform Society: Public Values In A Connective World.* Oxford University Press. https://doi.org/10.1093/oso/9780190889760.001.0001

Eger, M.A. & S. Valdez. (2014). Neo-nationalism in Western Europe. *European Sociological Review*, *31*:1, 115-130. https://doi.org/10.1093/esr/jcu087

Enli, G. & E. Skogerbø. (2013). Personalised Campaigns in Party-Centered Politics: Twitter and Facebook as Arenas for Political Communication. *Information Communication and Society, 16*(5), 757-774.

Elibol, R., S. Hagen & C. van de Ven. (2019). De krant voor de Carmiggelt-man. *De Groene Amsterdammer*. Retrieved from https://www.groene.nl/artikel/de-krant-voor-de-carmiggelt-man

Eriksson, M., et al. (2018). *Spotify Teardown: Inside the Black Box of Streaming Music*. MIT Press. https://doi.org/10.7551/mitpress/10932.001.0001

Es, K. Van, D. Van Geenen & T. Boeschoten. (2014). Mediating the Black Pete discussion on Facebook: Slacktivism, flaming wars, and deliberation. *First Monday, 19*(12). http://dx.doi.org/10.5210/fm.v19i12.5570

Estrada-Jiménez, J., J. Parra-Arnau, A. Rodríguez-Hoyos & J. Forné. (2019). On the regulation of personal data distribution in online advertising platforms. *Engineering Applications of Artificial Intelligence*, *82*: 13-29. https://doi.org/10.1016/j.engappai.2019.03.013

Feller, A., Kuhnert, M., Sprenger, T., & Welpe, I. (2011, July). *Divided They Tweet: The Network Structure of Political Microbloggers and Discussion Topics*. Paper presented at the Fifth International AAAI Conference on Weblogs and Social Media. Retrieved from http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/viewFile/2759/3210.

Fischer-Bollin, P et al. (2020). *Nationalismus in Europa - Einheit in Vielfalt?*. Retrieved from https://www.kas.de/de/einzeltitel/-/content/nationalismus-in-europa-einheit-in-vielfalt

Flaxman, S., S. Goel & J.M. Rao. (2016). Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly*, *80*, 298-320.

Fletcher, R. (2019). *The Rise of Populism and the Consequences for News and Media Use*. Retrieved from http://www.digitalnewsreport.org/survey/2019/the-rise-of-populism-and-the-consequences-for-news-and-media-use/

Friggeri, A., L.A. Adamic, D. Eckles & J. Cheng. (2014, May). *Rumor Cascades*. Paper presented at the Eighth AAAI Conference on Weblogs and Social Media. Retrieved from https://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/view/8122/8110

Fuchs, C. (2010). Alternative Media as Critical Media. *European Journal of Social Theory*, *13*(2), 173-192. https://doi.org/10.1177/1368431010362294

Fuchs, C. (2020). *Nationalism On The Internet: Critical Theory and Ideology in the Age of Social Media and Fake News*. Routledge.

Galpin, C., H. Trenz. (2017). The Spiral of Euroscepticism: Media Negativity, Framing and Opposition to the EU. In M. Caiani & S. Guerra (Eds.), *Euroscepticism, Democracy and the Media* (pp. 49-72). Palgrave Macmillan. https://doi.org/10.1057/978-1-137-59643-7_3

Geenen, van D. et al. (2016). Mining One Week Of Twitter. Mapping Networked Publics In The Dutch Twittersphere. *Selected Papers of AoIR 2016: The 17th Annual Conference of the Association of Internet Researchers*. Retrieved from https://dspace.library.uu.nl/bitstream/handle/1874/358313/1394_1362_1_PB.pdf?sequence=1

Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press.

Grover, T. & G. Mark. (2019, June). *Detecting Potential Warning Behaviors of Ideological Radicalization in an Alt-Right Subreddit*. Paper presented at the Thirteenth International AAAI Conference on Web and Social Media, Munich, Germany. https://www.aaai.org/ojs/index.php/ICWSM/article/view/3221

Guess, A., B. Nyhan, B. Lyons & J. Reifler. (2018). Avoiding The Echo Chamber About Echo Chambers: Why selective exposure to like-minded political news is less prevalent than you think (Knight foundation report). Retrieved from https://kf-site-production.s3.amazonaws.com/media_elements/files/000/000/133/original/Topos_KF_White-Paper_Nyhan_V1.pdf

Hall, S. (1973, September). *Encoding and decoding in the television discourse*. Paper presented at the Council of Europe Colloquy on Training in the Critical Reading of Televisual Language, Centre for Mass Communications Research, Leicester University.

Hess, D.J. & M. Renner. (2019). Conservative political parties and energy transitions in Europe: Opposition to climate mitigation policies. *Renewable and Sustainable Energy Reviews*, *104*, 419-428.

Heydari, A., J. Zhang, S. Appel, X. Wu & G. Ranade. (2019). *YouTube Chatter: Understanding Online Comments Discourse on Misinformative and Political YouTube Videos*. Retrieved from https://arxiv.org/abs/1907.00435

Hilhorst, S., & Hermes, J. (2016). "We have given up so much": Passion and denial in the Dutch Zwarte Piet (Black Pete) controversy. *European Journal of Cultural Studies*, *19*, 218–233. https://doi.org/10.1177/1367549415603381

Humphreys, A. (2015). *Social Media: Enduring Principles*. Oxford University Press.

Hussein, E., P. Juneja & T. Mitra. (2020, May). *Measuring Misinformation in Video Search Platforms: An Audit Study on YouTube*. Paper presented at the 23rd ACM Conference on Computer-Supported Cooperative Work and Social Computing. https://doi.org/10.1145/3392854

Hussain, M., S. Tokdemir, N. Agarwal & S. Al-Khateeb. (2018, August). Analyzing Disinformation and Crowd Manipulation Tactics on YouTube. Paper presented at the International Conference on Advances in Social Network Analysis and Mining, Barcelona, Spain. https://doi.org/10.1109/ASONAM.2018.8508766

Jacomy, M., T. Venturini, S. Heymann & M. Bastian. (2014). ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLOS ONE, 9*(6). https://doi.org/10.1371/journal.pone.0098679

Just, N. & M. Latzer. (2016). Governance by algorithms: reality construction by algorithmic selection on the Internet. *Media, Culture & Society, 39*(2), 238-258. https://doi.org/10.1177/0163443716643157

Kaiser, J. & A. Rauchfleisch. (2019). The implications of venturing down the rabbit hole. *Internet Policy Review*. Retrieved from https://policyreview.info/articles/news/implications-venturing-down-rabbit-hole/1406

Katz, E. & P.F. Lazarsfeld (1955). *Personal Influence: the part played by people in the flow of mass communications*. Free Press.

Kraker, J. De, S. Kuys, R.J.M. Cörvers & A. Offermans. (2014). Internet public opinion on climate change: a world views analysis of online reader comments. *International Journal of Climate Change Strategies and Management*, *6*(1), 19-33. https://doi.org/10.1108/IJCCSM-09-2013-0109

Klinger, U. & J. Svensson. (2018). The end of media logics? On algorithms and agency. *New Media & Society*, *20*(12), 4653-4670. https://doi.org/10.1177/1461444818779750

Kuiken, J., A. Schuth, M. Spitters & M. Marx. (2017). Effective Headlines of Newspaper Articles in a Digital Environment. *Digital Journalism*, *5*(10), 1300-1314. https://doi.org/10.1080/21670811.2017.1279978

Lerner, A., A.K. Simpson, T. Kohno & F. Roesner. (2016, August). *Internet Jones and the Raiders of the Lost Trackers: An Archaeological Study of Web Tracking from 1996 to 2016*. Paper presented at the 25th USENIX Security Symposium, Austin, TX. https://doi.org/10.5555/3241094.3241172

Manovich, L. (2009). *Cultural Analytics: Visualizing Cultural Patterns in the Era of "More Media"*. Retrieved from http://manovich.net/index.php/projects/cultural-analytics-visualizing-cultural-patterns

Manovich, L. (2016). The Science of Culture? Social Computing, Digital Humanities and Cultural Analytics. *Journal of Cultural Analytics*. https://doi.org/10.22148/16.004

Mansour-Ille, D. (2020). Social Media and the Dynamics of Radicalization and Violent Extremism among Female Migrant Workers. *Studies in Ethnicity and Nationalism*, *19*(3), 248-268. https://doi.org/10.1111/sena.12309

McCreery, M.P. & S.K. Krach. (2018). How the human is the catalyst: Personality, aggressive fantasy, and proactive-reactive aggression among users of social media. *Personality and Individual Differences*, *133*(15), 91-95. https://doi.org/10.1016/j.paid.2017.06.037

Mehrnezhad, M. (2020). *A Cross-Platform Evaluation of Privacy Notices and Tracking Practices*. Retrieved from https://eusec20.cs.uchicago.edu/eusec20-Mehrnezhad.pdf

Mertens, S., L. d'Haenens & R. De Cock. (2019). Online News Consumption and Public Sentiment toward Refugees: Is there a Filter Bubble at Play? Belgium, France, the Netherlands, and Sweden: A Comparison. In L. d'Haenens, W. Joris & F. Heinderyckx (Eds.), *Images of Immigrants and Refugees in Western Europe: Media Representations, Public Opinion, and Refugees' Experiences* (pp. 141-158). Leuven University Press.

O'Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Penguin Books.

Pariser, E. (2011). *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. Penguin.

Park, C.S. (2015). Applying "Negativity Bias" to Twitter: Negative News on Twitter, Emotions, and Political Learning. *Journal of Information Technology & Politics*, *12*(4), 342-359. https://doi.org/10.1080/19331681.2015.1100225

Postel-Vinay, K. (2017). *How neo-nationalism went global*. Retrieved from https://theconversation.com/how-neo-nationalism-went-global-74095

Rhodes, L. (Producer) & J. Orlowski (Director). (2020). *The Social Dilemma*. Exposure Labs.

Ribeiro, M.H. et al. (2020, January). *Auditing radicalization pathways on YouTube*. Paper presented at the 2020 Conference on Fairness, Accountability, and Transparency. https://doi.org/10.1145/3351095.3372879

Roberts, S. (2018). Digital detritus: 'Error' and the logic of opacity in social media content moderation. *First Monday*, *23*(3). http://dx.doi.org/10.5210/fm.v23i3.8283

Rogers, R. (2013). *Digital Methods*. MIT Press.

Rosenberg, D. (2013). Data before the Fact. In: L. Gitelman (Ed.), *"Raw Data" is an Oxymoron* (pp. 15-40). MIT Press.

Ruckenstein, M. & J. Granroth. (2020). Algorithms, advertising and the intimacy of surveillance. *Journal of Cultural Economy*. 13(1), 12-24. https://doi.org/10.1080/17530350.2019.1574866

Schäfer, M.T. (2019). *A Shattered Public Sphere. How Social Media Platforms Transform The Political Debate*. Retrieved from http://mtschaefer.net/media/uploads/docs/A-Shattered-Public-Sphere_publish-1.pdf

Scott, K. (2015). The Pragmatics of Hashtags: Inference and Conversational Style on Twitter. *Journal of Pragmatics*, *81*, 8-20. https://doi.org/10.1016/j.pragma.2015.03.015

Shen, C., et al. (2018). Fake images: The effects of source, intermediary, and digital media literacy on contextual assessment of image credibility online. *New Media & Society*, *21*(2), 438-463. https://doi.org/10.1177/1461444818799526

Shirky, C. (2011). The Political Power of Social Media: Technology, the Public Sphere, and Political Change. *Foreign Affairs*, *90*(1), 28-41. https://www.jstor.org/stable/25800379

Sieben, I. (2019). De stille meerderheid. *Mens en maatschappij, 93*(4), 309-311. https://doi.org/10.5117/MEM2018.4.001.SIEB

Smith, M, C. Zongott, B. Henn & G. Von Voigt. (2012, June). *Big data privacy issues in public social media*. Paper presented at the 6th IEEE International Conference on Digital Ecosystems and Technologies, Campione d'Italia, Italy. https://doi.org/10.1109/DEST.2012.6227909

Soffer, O. (2019). Algorithmic Personalization and the Two-Step Flow of Communication. *Communication Theory*. https://doi.org/10.1093/ct/qtz008

Spohr, D. (2017). Fake news and ideological polarization: Filter bubbles and selective exposure on social media. *Business Information Review*, *34*(3), 150-160. https://doi.org/10.1177/0266382117722446

Stieglitz, S. & L. Dang-Xuan. (2013). Social media and political communication: a social media analytics framework. *Social Network Analysis and Mining, 3*, 1277-1291. https://doi.org/10.1007/s13278-012-0079-3

Stier, S., A. Bleier, H. Lietz & M. Strohmaier. (2017). Election Campaigning on Social Media: Politicians, Audiences, and the Mediation of Political Communication on Facebook and Twitter. *Political Communication, 35*(1), 50-74. https://doi.org/10.1080/10584609.2017.1334728

Terrill, R.E. (2017). The Post-Racial and Post-Ethical Discourse of Donald J. Trump. *Rhetoric and Public Affairs, 20*(3), 493-510. https://doi.org/10.14321/rhetpublaffa.20.3.0493

Traag, V. A., L. Waltman & N. J. van Eck. (2019). From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, *9*(1). https://doi.org/10.1038/s41598-019-41695-z

Udupa, S. (2019). Nationalism in the Digital Age: Fun as a Metapractice of Extreme Speech. *International Journal of Communication*, *13*(1), 3143–3163. https://doi.org/10.5282/ubm/epub.69633

Vargo, C.J. & T. Hopp. (2020). Fear, Anger, and Political Advertisement Engagement: A Computational Case Study of Russian-Linked Facebook and Instagram Content. *Journalism & Mass Communication Quarterly*, *97*(3), 743-761. https://doi.org/10.1177/1077699020911884

Veerbeek,

Wieneke, A. & C. Lehrer. (2016). Generating and exploiting customer insights from social media data. *Electronic Markets*, *26*(3): 245–268. https://doi.org/10.1007/s12525-016-0226-1

Wieringa, M., T. de Winkel & C. Lewis. (2017). *Wie is de waakhond op sociale media?*. https://doi.org/10.13140/RG.2.2.31643.98084

Wieringa, M., D. van Geenen, M.T. Schäfer & L. Gorzeman. (2018). Political topic-communities and their framing practices in the Dutch Twittersphere. *Internet Policy Review*, *7*(2), 1-16. https://doi.org/10.14763/2018.2.793

Williams, H., J. McMurray, T. Kurz, F. Lambert. (2015). Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Global Environmental Change*, *32*, 126-138. https://doi.org/10.1016/j.gloenvcha.2015.03.006

Woolley, S. (2016). Automating power: Social bot interference in global politics. *First Monday*, *21*(4). https://doi.org/10.5210/fm.v21i4.6161

Zhao, X., et al. (2013, April). *The Many Faces of Facebook: Experiencing Social Media as Performance, Exhibition, and Personal Archive*. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/2470654.2470656

Zuiderveen Borgesius, F., et al. (2016). Should we worry about filter bubbles?. *Internet Policy Review*, *5*(1). https://doi.org/10.14763/2016.1.401